

**MULTI-WAVELENGTH CLOCK CONTROL UNIT MW-CCU:
SUMMARY REPORT**

ESA contract: 4000132316/20/NL/AR/zk

**IAN HILL
GEOFF BARWOOD
CONOR ROBINSON
PETER TSOULOS
JONATHAN STACEY
ALESSIO SPAMPINATO
PAUL GAYNOR
PATRICK GILL**

NPLML - COMMERCIAL

FEBRUARY 2022

**MULTI-WAVELENGTH CLOCK CONTROL UNIT:
Summary report**

ESA contract: 4000132316/20/NL/AR/zk

IAN HILL
GEOFF BARWOOD
CONOR ROBINSON
PETER TSOULOS
JONATHAN STACEY
ALESSIO SPAMPINATO
PAUL GAYNOR
PATRICK GILL

© NPL Management Limited, 2022

National Physical Laboratory
Hampton Road, Teddington, Middlesex, TW11 0LW

This report is NPL - Commercial and must not be exposed to casual examination. It is not for general distribution and should not be cited as a reference other than in accordance with the contract.

GLOSSARY/ABBREVIATIONS

ACES	Atomic clock ensemble in space
CTE	Coefficient of thermal expansion
DAC	Digital to analogue converter
DDS	Direct digital synthesis
EOM	Electro-optic modulator
FEA	Finite element analysis
HSL	High stability laser
ISS	International space station
LO	Local oscillator
MAIT	Manufacture, assembly, integration and test
PDH	Pound-Drever-Hall
PID	Proportional-integral-derivative
PLL	Phase locked loop
PMR	Project Management Review
SWaP	Size, weight and power
TRL	Technical readiness level
ULE	Ultra-low expansion

EXECUTIVE SUMMARY

The multi-wavelength clock control unit requirements outlined in the ITT Statement of Work refer to a particular sub-system unit of an optical clock architecture such as that of the Sr optical lattice clock. Such sub-systems include a highly-stable clock probe laser (optical local oscillator) stabilised to an ultra-stable optical cavity, an atomic physics package whereby atoms / ions are cooled, state-prepared and manipulated close to the quantum ground state of their motion via cooling & auxiliary lasers prior to clock laser probing of an ultra-narrow clock transition. A further sub-unit is a femtosecond comb providing the “clockwork” down-conversion from the optical to the microwave region. In particular, the multi-wavelength clock control unit (MW-CCU) provides means for frequency stabilisation of all the cooling & auxiliary lasers at appropriate levels of stability necessary for full functionality of the optical clock.

This MW-CCU activity comprised 5 specific design tasks, beginning with a critical review of the original CCU project results, identifying design deltas for improved functional performance, environmental resilience, robustness and reduced SWaP. In particular, an enhanced design to achieve the primary clock laser stabilisation on one CCU axis has been considered.

The 2nd activity concerned finite element analysis to inform these design updates in respect of enhanced cavity mounting and holding arrangements that will cope with vibration and shock conditions during launch and deployment, whilst also ensuring the cavity’s vibration insensitivity in the operational phase. A trade-off between cubic cavity size and the effect on cavity performance was also carried out. The 3rd task considered desirable opto-electronic and opto-mechanical upgrades to control differential drift between orthogonal axes in both thermal and isothermal cases. It also addressed the requirements for high-finesse low-thermal noise mirror coatings needed to stabilise the clock laser sub-harmonic at 1397 nm as well as the frequency doubling arrangement for direct CCU on-platform production of the 698 nm clock laser light. In respect of frequency control electronics, a fully digital design for the clock laser PDH control circuitry with few-MHz bandwidth will be built and tested, in order to meet the necessary control requirements in space operation.

The final activity involved generation of the manufacture, assembly, integration and test (MAIT) arrangements to implement the design upgrades identified. These will be carried out in a follow-on activity. To this end, a MAIT roadmap through to launch of a MW-CCU sub-system at EM level, via an in-orbit demonstration, was drawn up.

CONTENTS

1 INTRODUCTION.....	1
USING THE CCU FOR STRONTIUM ION CLOCKS AND RUBIDIUM SPECTROSCOPY	1
2 REVIEW OF ORIGINAL CCU PERFORMANCE.....	3
2.1 OPTICAL INPUT FOR THE CCU	4
2.2 CCU ELECTRONIC DESIGNS.....	5
2.3 MEASURED CAVITY FINESSES AT STRONTIUM LATTICE CLOCK WAVELENGTHS...8	8
2.4 THERMAL SHIELDING AND PRELIMINARY MEASUREMENTS OF THE RELATIVE DRIFT RATE OF THE CLOCK AXIS TO THE ORTHOGONAL DIRECTION	8
2.5 LASER LINEWIDTH AND FREQUENCY STABILITY MEASUREMENTS AT PTB.....	10
2.5.1 689 nm laser linewidth.....	10
2.5.2 813 nm laser linewidth.....	11
2.6 MEASUREMENTS BEFORE AND AFTER PROTON BOMBARDMENT AT BIRMINGHAM UNIVERSITY.....	11
3 CCU DESIGN IMPROVEMENTS.....	12
3.1 OPTICAL UPGRADE FOR THE CLOCK AXIS	12
3.2 DESIGN UPGRADES FOR REQUIRED ENVIRONMENTAL TESTING.....	14
3.3 IMPROVED CHARACTERISATION OF THE SPACER THERMAL EXPANSION.....	15
3.4 REDUCING THE SWAP FOOTPRINT	16
4 INTEGRATING THE CLOCK STABILISATION SYSTEM INTO THE CCU	17
4.1 ELECTRONICS UPGRADES	17
5 FINITE ELEMENT ANALYSIS OF CCU CAVITY DESIGN & MOUNTING.....	17
6 CUBE SIZE TRADE-OFF.....	43
6.2 MESH REFINEMENTS CHECKS	43
6.3 ANNULI GEOMETRY	46
6.4 3-CM CUBE (1" mirror with annuli).....	47
6.5 4-CM CUBE (1" mirror with annuli).....	48
6.6 3-CM CUBE AND 4-CM CUBE (both with ½" mirrors)	50
6.7 EFFECTIVE CTE.....	52
6.7.1 3-CM CUBIC CAVITY	52
6.7.2 4-CM CUBIC CAVITY	53
6.8 SUMMARY OF SECTIONS 5 AND 6	53
7 CCU DUAL AXIS CCU OPTO-ELECTRONIC PERFORMANCE.....	55
7.1 RECOMMENDATIONS FOR FUTURE DEVELOPMENT FROM THE FIRST PROJECT	56
7.2 CCU CHARACTERISATION	57
7.2.1 CAVITY FINESSE MEASUREMENTS	57

7.2.2 INITIAL CHARACTERISATION OF FREQUENCY STABILITY	57
7.2.3 SUMMARY OF PREVIOUS LINEWIDTH AND STABILITY MEASUREMENTS.....	59
7.3 IMPROVED CHARACTERISATION OF THE CCU DRIFT RATES	61
7.4 IMPROVED CHARACTERISATION OF THE CCU VIBRATION SENSITIVITY	65
7.5 DESIGN DELTAS ON THE CCU CLOCK AXIS	65
7.5.1 CAVITY MIRROR OPTIONS	65
7.5.2 LINEAR ABSORPTION LOSSES AT THE SUB-HARMONIC CLOCK WAVELENGTHS	66
7.5.3 FREQUENCY DOUBLING FROM 1397 NM TO 698 NM	67
7.6 DESIGN DELTAS FOR THE INPUT OPTICS.....	68
7.6.1 Mode Matching into the cavity	68
7.6.2 COTS COMPONENTS	70
7.6.3 COMPACT OPTICAL ISOLATORS.....	70
PHASE NOISE CANCELLATION	71
8 ELECTRONICS ENHANCEMENTS.....	72
8.1 ELECTRONICS SYSTEM ARCHITECTURE.....	73
8.1.1 POWER SUPPLIES	74
8.1.2 MICROCONTROLLER.....	75
8.1.3 COMMUNICATION.....	76
8.2 ELECTRONICS HARDWARE DESIGN – CHANNEL CARD	76
8.2.1 FPGA AND MICROPROCESSOR.....	77
8.2.2 MEMORY DEVICES.....	77
8.2.3 CHANNEL CARDS POWER SUPPLIES	78
8.2.4 CLOCK GENERATION AND DISTRIBUTION.....	79
8.2.5 ADCS AND DACS	79
8.2.5.1 ADCs.....	79
8.2.5.2 DACs.....	80
8.3 SOFTWARE AND FPGA FIRMWARE.....	81
8.3.1 FRONT-END INTERFACE.....	81
8.3.2 BACKPLANE SOFTWARE	84
8.3.3 CHANNEL CARD SOFTWARE.....	84
8.3.4 COMMUNICATION PROTOCOL FOR PARAMETER CONFIGURATION	84
8.3.5 CHANNEL CARD CONFIGURATION PARAMETERS	86
8.4 FPGA FIRMWARE.....	87
8.4.1 TOP-LEVEL DESIGN	87
8.4.2 DIGITAL FILTERS.....	89
8.5 VERIFICATION OF PDH LOCK	91
8.6 LOCKING SCHEME – FPGA SIMULATION	91

8.6.1 DEMONSTRATION SET-UP.....	91
8.6.2 RESULTS DISCUSSION.....	92
8.7 LOCKING SCHEME – ELECTRONICS IMPLEMENTATION.....	94
8.7.1 DEMONSTRATION SET-UP.....	94
8.7.2 RESULTS DISCUSSION.....	95
8.8 EVALUATION OF THE SERVO USING A 1064 NM DFB LASER.....	96
9 MANUFACTURING PLAN FOR AN ENHANCED CCU.....	97
9.1 INTRODUCTION	97
9.2 COMPACT CUBE OPTIONS.....	98
9.3 IMPROVED CAVITY ALIGNMENT PROCEDURE	98
9.4 PLANNED CMM MEASUREMENTS BY NPL DIMENSIONAL TEAM	100
9.5 SELECTION OF ULE ORIENTATION.....	100
9.6 UPGRADING THE CCU CLOCK AXIS	100
10 IMPLEMENTATION PLAN.....	101
TASK 1: CUBIC CAVITY & MOUNTING REFINEMENT STUDIES & DEVELOPMENT.....	101
TASK 2: VACUUM HOUSING REFINEMENT.....	101
TASK 3: OPTICS REFINEMENT FOR COMPLETE CLOCK AND AUXILIARY INSENSITIVE SYSTEM.....	102
TASK 4: OPTO-ELECTRONICS & OPTICS	102
TASK 5: SERVO-ELECTRONICS DESIGN REFINEMENT.....	102
TASK 6: INTEGRATION AND FULL UNIT PERFORMANCE TEST & ENVIRONMENTAL TEST	103
TASK 7: ADVANCE TO TRL7 VIA ENVIRONMENTAL TESTING PROGRAMME FOR FREE FLYER.....	103
10.1 CURRENT STATUS WITHIN DIFFERENT ESA PROJECTS	103
10.2 DEVELOPMENT ROADMAP	103
10.3 BUDGET ESTIMATE FOR PROPOSED ACTIVITIES	104

1 INTRODUCTION

This Summary Report of the Multi-Wavelength Clock Control Unit (MW-CCU) project carried out under ESA contract 4000132316/20. It comprises a review of the Clock Control Unit (CCU) developed within ESA contract 4000117290/16 for stabilising all the cooling and auxiliary lasers in a strontium lattice clock, and reports on the follow-on design upgrades for the cubic cavity mechanical design, opto-electronics and electronics control. The cavity is also able to stabilise the sub-harmonic (at 844 nm) of the strontium ion clock cooling laser at 422 nm.

The CCU comprises a dual axis cubic cavity (DACC) based on NPL's patented force-insensitive, rigidly mounted cubic cavity [1]. The DACC is mounted in an aluminium frame (Figure 1) within a small vacuum housing.

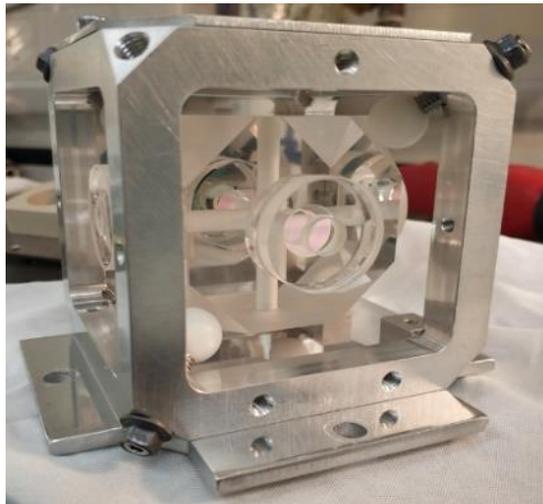


Figure 1: Mounted dual axis cubic cavity from ESA TRP CCU project

All the wavelengths required except for the clock wavelength at 698 nm are multiplexed down one axis of the DACC (Figure 4). For the Sr lattice clock, in addition to the clock laser wavelength of 698 nm, the auxiliary laser wavelengths required include:

- 461 nm (for 1st stage cooling of the atoms to low mK temperatures)
- 689 nm (for 2nd stage cooling to low μ K temperatures needed for lattice loading)
- 813 nm (for providing optical lattice confinement at the “magic” wavelength)
- 679 nm & 707 nm (for state repumping)

A critical step in the success of the DACC arrangement was the procurement of cavity mirrors with specified finesses covering the wavelength region from 461 nm to 922 nm. The fused silica substrates were 25.4 mm diameter, 6.35 mm thickness, 1 m radius of curvature with a super-polished finish on the concave side. The reverse side has an antireflection coating with a broadband reflectivity of $R < 0.25\%$ across the wavelength range of interest. In this TN, we summarise the results obtained to date with the CCU and then review the upgrade plans for the next stage of development.

1.1 USING THE CCU FOR STRONTIUM ION CLOCKS AND RUBIDIUM SPECTROSCOPY

Whilst the primary purpose of the clock control unit (CCU) relates to the frequency control of the cooling, auxiliary and (eventually) the clock laser for a strontium lattice clock [2, 3], the CCU was also designed to deal with two further scenarios:

- Frequency stabilisation of lasers to the rubidium D lines at 780 nm and 795 nm

- Frequency stabilisation of some of the lasers required for a trapped ion optical clock based on a single trapped $^{88}\text{Sr}^+$ ion.

The mirror reflectivity curve for the mirrors designed for CCU by FiveNine optics is shown in Figure 2. This shows a specified reflectivity value in the ~780 nm – 795 nm region that covers the first requirement. One of the CCU input ports (which normally would deal with light at 813 nm) could alternatively accept input at 780 nm or 795 nm.

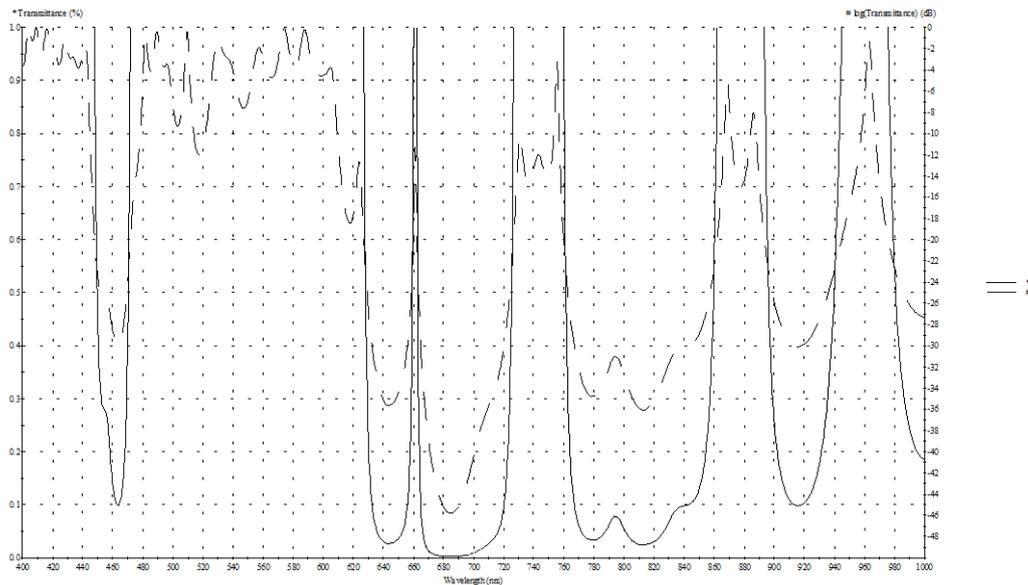


Figure 2 Broadband low-loss mirror coating transmittance vs wavelength in linear and log scale as provided by FiveNine Optics

The requirements for a $^{88}\text{Sr}^+$ optical clock [4] can be explained by reference to Figure 3. For this system, the clock frequency is the quadrupole transition at 674 nm; ultra-high finesse cavity mirrors with a dielectric coating that cover both 674 nm and 698 nm (the strontium lattice clock wavelength) can be commercially sourced. Light at the 422-nm cooling wavelength is typically provided by frequency doubling 844 nm. Light at 461 nm (required for photo-ionisation) could be generated either by frequency doubling 922 nm or provided directly from a bespoke GaN extended cavity diode laser. As shown in Figure 2, the mirrors in the CCU cover 844 nm, 922 nm and 461 nm. For optical clocks, we note that the frequency stability requirements for 922 nm and 844 nm will be similar, both being the sub-harmonic of cooling laser transitions with low tens MHz transition linewidths.

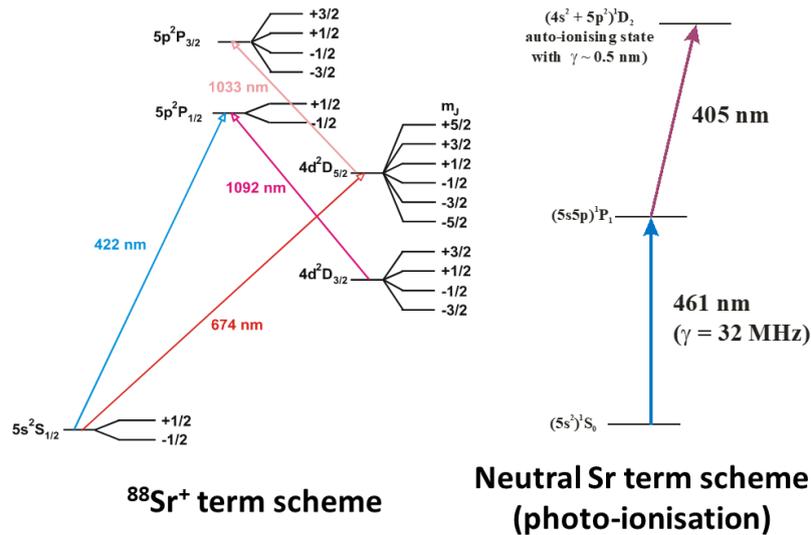


Figure 3: Term scheme for a $^{88}\text{Sr}^+$ ion clock (left) including the lasers (right) required for photo-ionisation

Of the remaining wavelengths in a $^{88}\text{Sr}^+$ clock, the light at 405 nm is provided by a multi-mode GaN laser and therefore does not require frequency stabilisation. However, an initial check with a wavemeter should be undertaken to set the laser drive current and operating temperature. The 405 nm laser drives Sr to an auto-ionising state that has a ~ 0.5 nm linewidth; providing this laser is particularly straightforward as this wavelength is that required for high-definition video recording. The clear-out lasers at 1092 nm and 1033 nm do not require cavity stabilisation; at NPL and frequency stabilisation can be provided via a wavemeter [5]. An alternative approach is to use spontaneous emission sources at these wavelengths [6]. Therefore, a CCU could be used to frequency stabilise all the laser sources that require cavity stabilisation for $^{88}\text{Sr}^+$ giving this device multiple applications in the development of optical clocks in space.

2 REVIEW OF ORIGINAL CCU PERFORMANCE

In this section, we review the performance of the CCU, with specific reference to Annex B of the statement of work. Some of the key CCU requirements are:

- The CCU shall be capable of robust absolute frequency stabilization at 461 nm, 689 nm and 813 nm with an absolute deviation of less than 1 MHz, 1 kHz, and 100 kHz respectively over 1 h. Re-pumpers at 707 nm and 679 nm shall also be frequency stabilized by using the CCU (R-01). Any long-term drift should be taken into account to reach the specified optical frequencies for a duration of at least 3 years (R-02).
- The CCU shall provide an error signal with sufficient bandwidth and SNR to frequency lock lasers to the following linewidth (95% of the power shall be within a spectral bandwidth equal to the linewidth specified above – R-03)
 - 1 MHz for 461 nm
 - 100 kHz for 813 nm
 - 1 kHz for 689 nm
- Mass and volume guidelines (Values apply to the optics package. Electronic drivers, frequency controllers, etc. are excluded; R-07):
 - CCU mass < 15 kg

- CCU volume < 35 l
- For the 689 nm laser the spectral power density in a frequency range of 1 kHz – 100 kHz away from the carrier shall be -65 dBc/Hz (this is R-04).
- R-05 states that It shall be possible to tune the frequency of the lasers to ± 10 MHz within 2 s. As frequency tuning is via tuning of an EOM drive frequency [7], we do not anticipate this to be a problem.
- R-06 states: The CCU shall allow for frequency modulation of the 689 nm laser at a modulation frequency of 30 – 50 kHz with a modulation amplitude of 2 MHz peak-to-peak relative to the given centre frequency. It shall be possible to switch on and off the frequency modulation via a TTL input. Modulation frequency and amplitude shall be adjustable. It shall be possible to connect different centre frequencies by TTL control for switching the modulation. This was not implemented in the existing CCU but can be implemented in a future electronics design.
- R-07: Mass and volume guidelines (note: values apply to the optics package. Electronic drivers, frequency controllers, etc. are excluded):
 - CCU mass < 15 kg,
 - CCU volume < 35 l

We achieve a total mass of 9912 g for the opto-mechanical assembly and enclosure and 6213 g for the dual axis cubic cavity and vacuum housing, giving a combined total of 16.125 kg. This exceeds the original requirement by 1.125 kg; however, we justify this by the decision made at the PDR to use a more complex cavity arrangement which provides the opportunity to stabilise the clock laser in a future iteration. Refinements to the vacuum housing and baseplate and enclosure could easily bring this into specification. The volume is 456 x 306 x 166 mm, or 23.2 litres, and is well within specification.

2.1 OPTICAL INPUT FOR THE CCU

The schematic for multiplexing the different wavelengths into the CCU is shown in Figure 4. Each of the wavelengths (i.e. 679 nm, 707 nm, 689 nm, 922 nm and either 813 nm, 844 nm, 780 nm or 795 nm) has a separate phase modulator (EOM) that can be driven at a different frequency. The dual-axis cubic cavity housed within a vacuum housing, together with the input fibre optics was packaged into a 19" rack unit (see photos in Figure 5 and Figure 6).

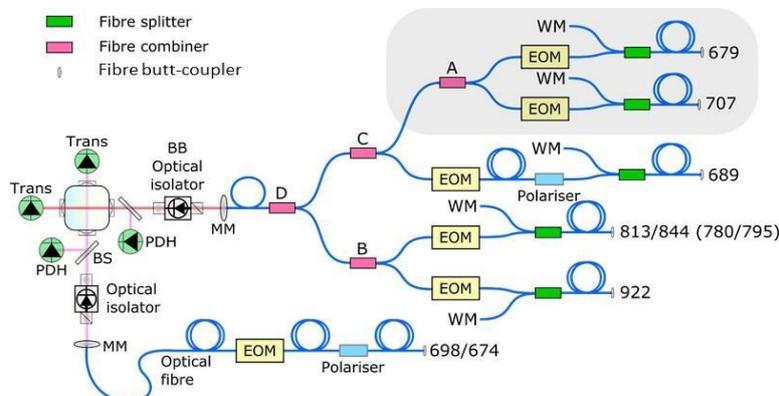


Figure 4: Schematic of the multi-wavelength dual-axis cubic cavity CCU DACC optical system. WM labels are for fibre to be routed to a multi-channel fibre switcher input to a commercial High-Finesse wavemeter. The shaded region around the 679 and 707 nm repump inputs indicates the option to omit

these wavelengths from the cavity stabilisation aspect of the CCU, as suitable stabilisation may be achieved using a High Finesse wavemeter.

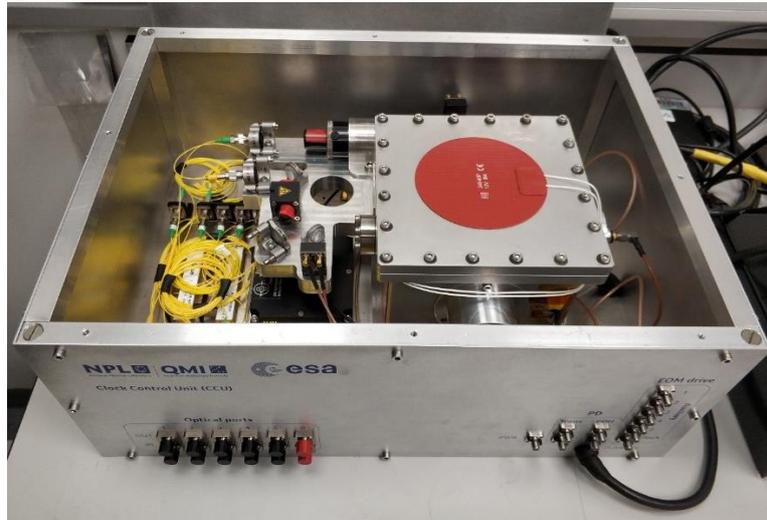


Figure 5: View of the CCU cavity vacuum chamber and fibre and free-space optics

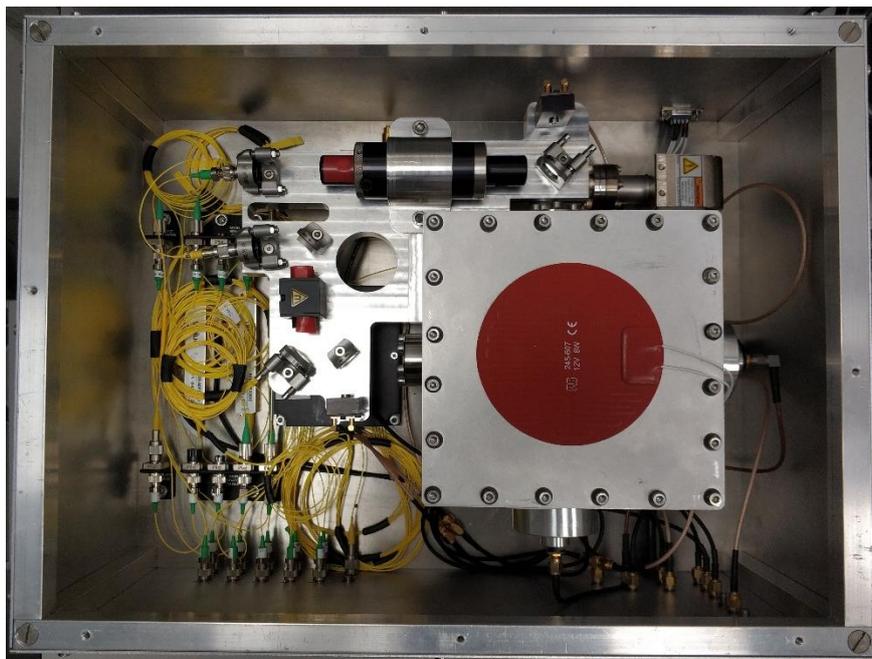


Figure 6: Top view looking down on the CCU optics package

2.2 CCU ELECTRONIC DESIGNS

The schematic for the electronics arrangement is shown in Figure 7.

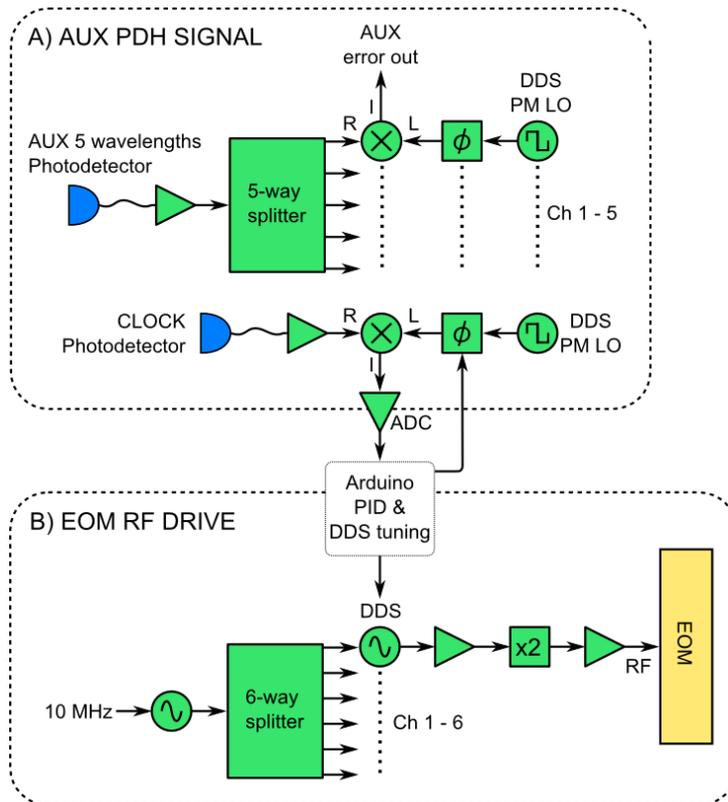


Figure 7: Functional overview of the CCU electronics module

The electronics are implemented partly via discrete Mini-Circuits components and partly via NPL-designed digital hardware. The 19" rack module comprising the electronics is shown in Figure 8 and Figure 9.



Figure 8: CCU electronics module

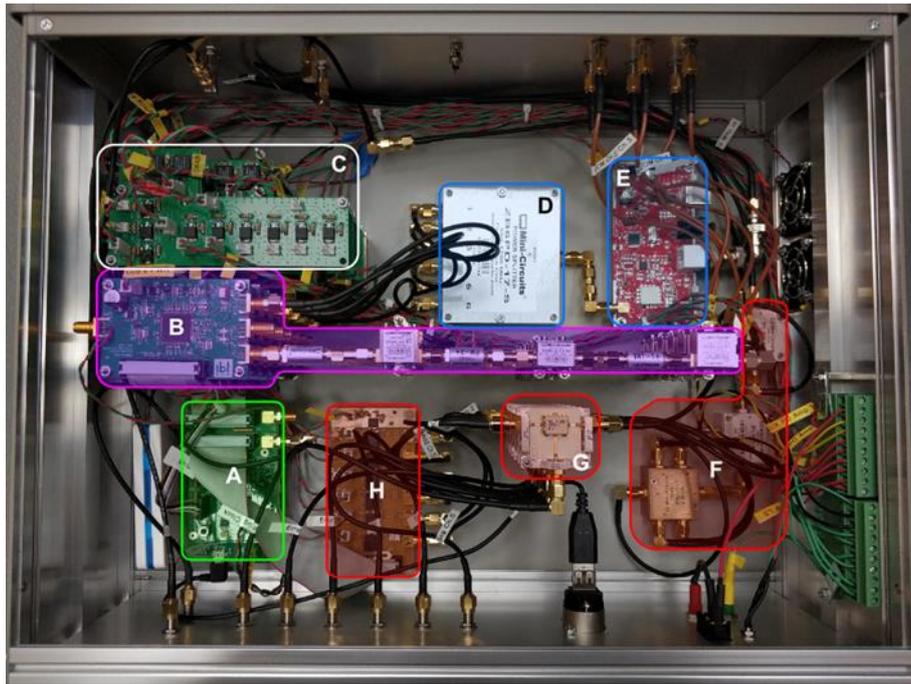


Figure 9: CCU electronics module. Electronic sub-systems are highlighted for clarity. A is the Arduino, B DDS and signal processing, C power regulation, D 6-way distribution of the DDS master clock, E 10 MHz PLL and 1 GHz DDS clock generation, F AUX and CLOCK PDH PD amplification and AUX 5-way splitting, G 6 PDH phase detectors, and H 6 PDH LO phase delay chips

Pound-Drever-Hall (PDH) error signals are recovered from a single photodiode at distinct modulation frequencies; the output from the detector in reflection is split several ways, with each split output separately demodulated in an analogue mixer. In this way, separate PDH error signals were recovered for each input wavelength. The modulation frequencies were selected to be sufficiently different so as to not limit the servo bandwidth while avoiding interference between channels from harmonics or intermodulation; frequencies of 8.900 MHz, 6.940 MHz, 16.625 MHz and 12.500 MHz were used for the locks at 689 nm, 698 nm, 813 nm and 922 nm respectively. A system evaluation demonstrated that there was minimal crosstalk between these different wavelength channels. Python software was developed to change different input parameters in the electronics including proportional-integral-differential (PID) gain parameters and the servo sense (polarity); engaging the cavity lock was also settable via python code.

The experimental PDH signals are shown in Figure 10 at each of the four different wavelengths. The PDH profiles are all different in shape as a result of both the different cavity finesses and modulation frequencies for each wavelength. The centre frequencies of each of these wavelengths are determined by the cavity resonant frequencies. We need a way to bridge the frequency interval between the cavity mode frequency and that of the strontium transition. For this, we have implemented the electronic sideband version of the double modulation method in [7]. Here, the modulator is driven with a phase modulated RF carrier; this carrier is typically in the hundreds of MHz to low GHz region and is used to bridge the atom to cavity frequency interval. The phase modulation corresponds to the PDH modulation.

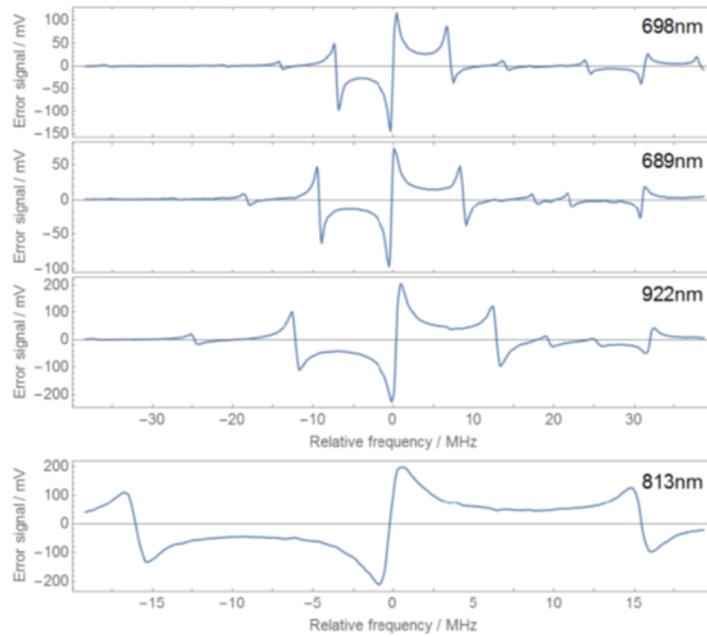


Figure 10: Simultaneously recorded PDH error signals from the dual axis cubic cavity ESA TRP CCU project. Note that the 698 nm clock PDH signal would be replaced by a 1397 nm PDH signal, which would be much narrower in MW-CCU case due to the very high finesse coating used on the clock axis

2.3 MEASURED CAVITY FINESSES AT STRONTIUM LATTICE CLOCK WAVELENGTHS

The cavity finesse at the different wavelengths in a strontium lattice clock were measured for the CCU and the results are shown in Table 1.

Table 1: Measured cavity finesse and inferred cavity mirror loss using specified mirror reflectance.

Measurement date	Axis	Wavelength (nm)	Fringe FWHM (kHz)	Measured finesse	Expected finesse (20 ppm loss)	~ Inferred loss (ppm)
9/3/2018	Clock	698.4	266	11.3 k	26 - 38 k	< 200
8/5/2018	Aux	689.5	342	8.8k	52 - 74 k	250 - 290
8/5/2018	Aux	813.4	1200	2.5 k	7 - 14 k	820 - 1050
8/5/2018	Aux	921.7	811	3.7 k	2.2 - 3.8 k	< 5

The orthogonal axis of the DACC is for the clock wavelength (698 nm for a strontium lattice clock, or 674 nm for Sr⁺) and used for monitoring length changes in the cavity spacer by reference to a stable clock laser.

2.4 THERMAL SHIELDING AND PRELIMINARY MEASUREMENTS OF THE RELATIVE DRIFT RATE OF THE CLOCK AXIS TO THE ORTHOGONAL DIRECTION

The cubic cavity was mounted within an aluminium frame which, in turn, had a single-layer heatshield attached. The complete cavity arrangement was then mounted within an aluminium

vacuum chamber. The thermal time constant was measured by applying a temperature step change and this was determined to be ~8.5 hours (Figure 11).

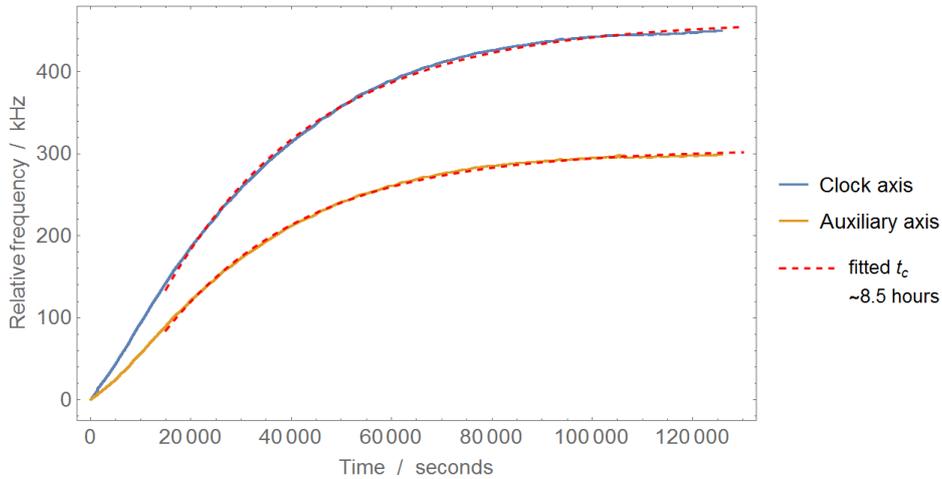


Figure 11: Measured frequency drift of the two cavity axes following a 300mK temperature step. The fit indicates a thermal response time for the cavity of about 8.5 hours

The temperature where the coefficient of thermal expansion is zero was determined via a temperature sweep and gave results of 24°C and 26°C for the auxiliary and clock axes respectively indicating that the ULE is not isotropic. It will therefore not be possible to control the cavity at a temperature where the thermal expansion is zero on both axes simultaneously. Also, from our modelling of the effect of using fused silica mirrors on a ULE spacer with thermal compensating rings [8] for our specific geometry, we conclude that the ULE purchased had an intrinsic CTE crossover temperature at the upper end of the specified 25°C – 35°C region.

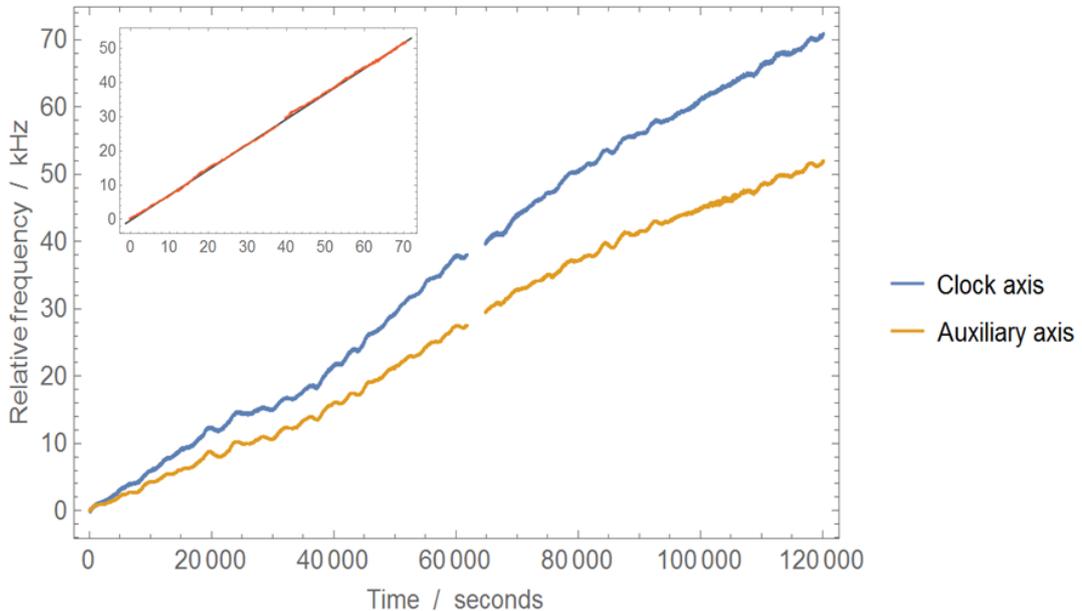


Figure 12 Preliminary measurement of the isothermal frequency drift of the two cavity axes. Inset: plotted against each other, the drift of the two axes of the cavity are highly correlated.

The relationship between the thermal expansion of the clock axis and the orthogonal direction was measured as part of the first CCU project (Figure 12) and further measurements are planned within this activity to assess the stability of the relation. In a future system, we plan that the DACC will also provide the short-term reference for the clock laser, offering an integrated clock and auxiliary laser stabilisation unit with significantly reduced overall system

complexity. Using the measured physical relation between length changes in each orthogonal cavity axis, the drift of the cavity at the other wavelengths can be removed by an electronic feed-forward scheme.

The DACC spacer ULE material is specified to have a zero CTE within the temperature range 25°C to 35°C. However, having fused silica mirror substrates (which leads to a lower thermal noise limit) means that the effective zero CTE temperature is decreased close to 0 °C, significantly below that for an all-ULE spacer and mirror substrate arrangement. The use of ULE compensating rings goes some way [8] towards pulling back the zero CTE towards the all-ULE temperature. However, the effective zero CTE temperature is still reduced by several degrees for our configuration of a 5-cm spacer and 25.4 mm diameter mirrors. Nevertheless, the effective zero CTE temperature for an NPL cubic cavity can still be above room temperature [9], if the ULE crossover temperature for the material supplied is close to 35°C, as is observed from the data above.

2.5 LASER LINEWIDTH AND FREQUENCY STABILITY MEASUREMENTS AT PTB

Characterisation of the clock control unit (CCU) cavity was undertaken via a beat frequency comparison at selected wavelengths against independently stabilised lasers both at NPL and the Physikalisch-Technische Bundesanstalt (PTB) in Germany. Measurements using a selected sub-set of the full laser set were used to infer operation within specification over the full wavelength region. Testing of the CCU cavity was undertaken at PTB using their laser sources at 689 nm and 813 nm; a clock laser at 698 nm was also used. This laser was stabilised to a separate ULE cavity [10] that was, in turn, locked to a cryogenic silicon cavity [11]. At 813 nm, the beat linewidth was 200 kHz and a noise level of -60 dBc/Hz was measured at an offset of 300 kHz. The frequency drift over one hour at this wavelength was below 4 kHz, well within the 100 kHz requirement. At 689 nm, a maximum frequency deviation of 0.5 kHz/hour was measured, within the 1 kHz/hour requirement, using the measured relative drift between the clock and auxiliary axis. A drift rate suppression factor of 7.2 was demonstrated at 689 nm.

2.5.1 689 NM LASER LINEWIDTH

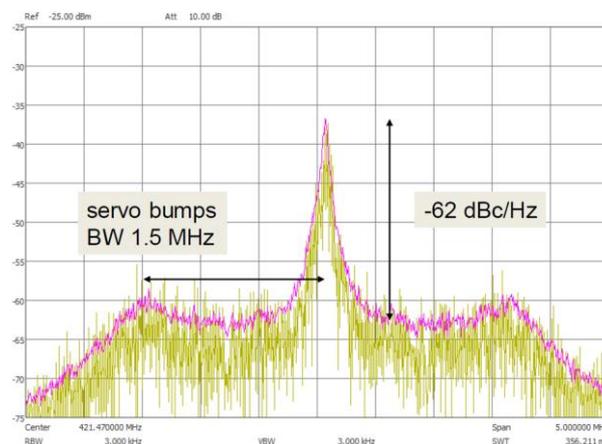


Figure 13: Beat of the 689 nm ECDL with Sr 689 nm cooling laser

The linewidth of the 689 nm laser was measured at PTB by beat frequency versus the laser used for their strontium lattice clock. The 689 nm beat signal could be also counted by a zero dead-time frequency counter and analysed with a RF spectrum analyser (Figure 13 and Figure 14). The position of the servo bumps indicates a bandwidth of the servo loop of about 1.5 MHz.

From the beat spectra, a power spectral density of the beat of -62 dBc/Hz at 250 kHz offset can be read off. From the error signal, no excess noise at lower frequencies was observed. As the noise floor in the spectrum contains both lasers, assuming similar behaviour, the requirement of -65 dBc/Hz is fulfilled. The 3 dB linewidth of the beat signal is about 3 kHz, limited by the resolution bandwidth of the spectrum analyser.

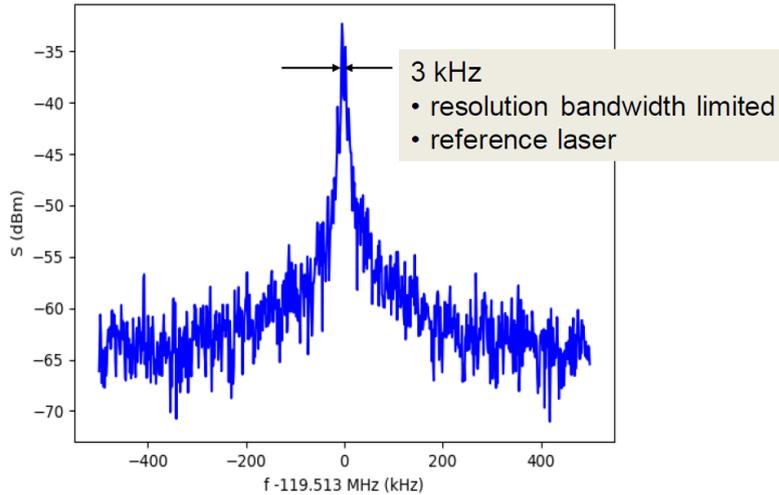


Figure 14: Central part of the beat showing 3 kHz resolution-limited linewidth

2.5.2 813 NM LASER LINEWIDTH

Using a similar beat frequency comparison arrangement to that at 689 nm, the 813 nm laser was locked to the CCU. Here the beat spectrum showed a noise of -60 dBc at an offset of 300 kHz (Figure 15).

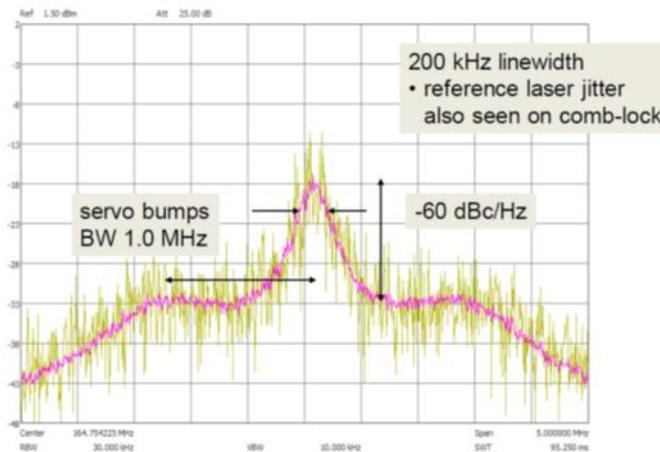


Figure 15: Beat of 813 nm laser to Sr lattice laser

2.6 MEASUREMENTS BEFORE AND AFTER PROTON BOMBARDMENT AT BIRMINGHAM UNIVERSITY

The final step in the testing of the CCU cavity involved a test of the cavity resilience to irradiation by proton bombardment (Figure 16). The effect on the cavity mirrors was determined by measurements of the cavity finesse before and after irradiation. The finesse was determined at 698 nm by scanning a narrow linewidth laser over the TEM₀₀ mode of the cubic cavity. The cubic cavity, in its vacuum housing, was irradiated on a 36 MeV proton beam line at the University of Birmingham to a total fluence of 4×10^{10} protons·cm⁻², targeted along the cavity axis, centred on the mirror. Measurements of the finesse were 11,270 and 11,240 before and

after irradiation respectively. These two results are consistent and indicate no perceptible performance degradation resulting from the experienced radiation dose.

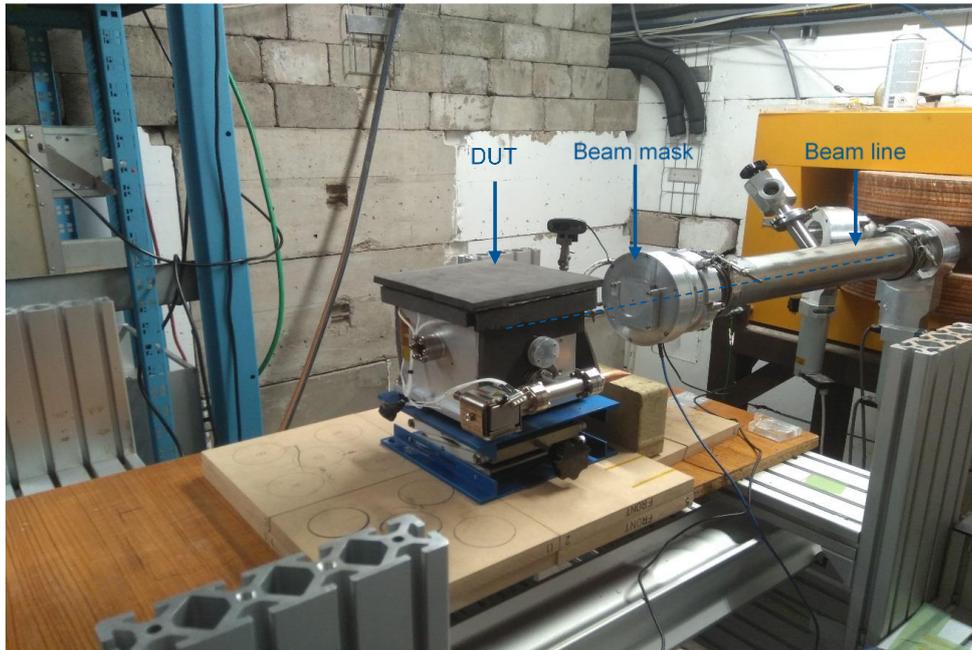


Figure 16: General experimental arrangement on the Birmingham beamline. The CCU vacuum housing can be seen in the centre of this photo

3 CCU DESIGN IMPROVEMENTS

The results of the first study highlight the successful development of a compact clock control unit for a future space strontium lattice clock [12]. The CCU has also been designed for use with a strontium ion clock [4] and clear out radiation provided either via wavemeter controlled lasers [5] or spontaneous emission sources [6]. A critical requirement for future activity is to increase the technology readiness level (TRL) of the cubic cavity and mounting arrangement. However further measurements and design changes will clearly also enhance the laboratory device. For example, by replacing the clock cavity axis mirrors with high-finesse alternatives, the same cavity could provide the reference for all the required lasers, including 698 nm. For a fully optimised system, crystalline mirror coatings [13] can be expected to reduce the thermal noise by more than a factor of two below that of dielectric mirror coating. However, these mirrors are currently restricted to longer-wavelength sources and do not work well at 698 nm. In this case, therefore, it would be better to use a laser at 1397 nm and generate the clock light by frequency doubling. Using high-finesse mirrors in the dual-axis cubic cavity at the clock wavelength would also allow verification of the low sensitivity of the cube to vibration. From modelling and previous experimental results, we anticipate that this will be at the low parts in $10^{-11}/g$ [1]. This cavity, together with crystalline mirrors in the clock axis and some additional vibration isolation arrangements, can therefore be expected to achieve parts in 10^{16} frequency instability over timescales from 0.1 s to 100 s, as necessary for lattice clocks operating with instability $<1 \times 10^{-15}$ per square root of averaging time

3.1 OPTICAL UPGRADE FOR THE CLOCK AXIS

To fully optimise the cavity performance as a clock laser reference, we also anticipate that further heat-shielding will be required to improved short-term temperature control. We will also improve on the measurements of the temperatures where the CTE is zero for the two cavity arms in order to optimise the set temperature for the clock axis.

In the following sections, we detail the proposed programme of work within this project to improve the performance of the CCU. These improvements can be summarised as follows:

- Improved characterisation of the differential drift between orthogonal axes of the cubic cavity (both thermal and isothermal). We will also investigate further the ULE anisotropy, particularly with regards to the different thermal expansion coefficients in different directions and make measurements over longer timescales. This will lead to improved differential drift compensation for the auxiliary laser wavelengths at 461 nm, 689 nm, 813 nm, 679 nm and 707 nm.
- Outline of the planned design upgrade needed on the clock cavity axis to include very-high-finesse low-thermal-noise crystalline mirror coating mirrors for operation at 1397 nm. A further upgrade to the CCU platform will also be designed in order to include the 1397 nm laser plus frequency doubler unit, thus providing the cavity-stabilised 698 nm clock laser light directly on-board the CCU. These modifications will include updating the input optics for 1397 nm (the clock laser sub-harmonic) rather than 698 nm. Detector upgrades are also required to implement InGaAs technology rather than Si avalanche photodiodes at this longer wavelength. Additional optics will be required for coupling into the frequency doubler.
- Finite element analysis of the dual axis 5-cm cubic cavity within the mounting frame to ensure maintenance of cavity alignment within the frame during the high vibration and shock environments encountered during launch and satellite deployment. Here, it is essential to avoid potential loss of frictional hold at the support points during this phase, which could result in small cavity rotations and resulting loss of optimal vibration insensitivity in the post-launch operational phase. It is intended to investigate an improved re-location strategy via a spacer upgrade incorporating cone-vee-flat location on the vertex support points.
- Investigation, using finite element analysis, of the potential for size reduction of the cubic spacer to 3 cm, in order to assess e.g. how this affects the vertex cut depth and conduct a trade-off analysis against the impact on resulting vibration insensitivity

A key driver for the required improvements to the CCU design relates to the practical implementation of a high finesse cavity bore for stabilisation of light at the strontium lattice clock frequency (698 nm). The lowest thermal noise cavities leading to the best short-term laser frequency stabilities are based on crystalline mirrors [13]. However, crystalline mirror coatings have not been achieved directly at the 698 nm clock laser wavelength. As a result, our approach will be to make use of available crystalline mirror coating technology at 1397 nm for the clock laser cubic cavity axis (as has already been demonstrated at 1064 nm [4]) and frequency double to 698 nm to provide the clock laser output to probe the cold atoms. Crystalline mirror coatings are available with a reflectivity able to provide a cavity finesse of ~200,000 and a thermal noise limit projected to be in the mid parts in 10^{16} region for our 5-cm cubic cavity. Upgrading to crystalline mirrors at 1397 nm will have a number of consequences for the design and characterisation of the DACC, which we will detail further in this section:

- Development of a compact 1397 nm laser source, together with a frequency doubler to obtain 698 nm. A bespoke compact laser diode system is required, and one option is based on volume holographic grating technology commercially exploited by Sacher [14]. The clock axis input and detection optics will require some redesign to deal with 1397 nm rather than 698 nm. For example, silicon photodiodes will need replacing with InGaAs devices.

- Careful mounting of the cavity spacer within the mounting frame to guarantee minimal vibration insensitivity. A low level of sensitivity to vibration was not a critical requirement for the original CCU which did not provide the reference for the clock laser. The calculated cut-out depth at the vertex supports that gives this minimum sensitivity will be the same for both the new and the original CCU.
- Higher bandwidth servos that will be required for frequency stabilisation of lasers at 1397 nm that are expected to have larger linewidths than, for example, Nd-YAG lasers at 1064 nm which is where much of the ESA-funded NPL activity has centred to date.

Separate from the design upgrades required to use the DACC for clock laser frequency stabilisation, we need to address some further – design upgrade issues, as follows:

- Additional damping or feed-forward acceleration compensation techniques between the cubic cavity chamber and platform will be reviewed. MW-CCU will undertake design deltas in respect of the cubic cavity vertex mounting and holding arrangements, such as e.g. cone-vee-flat locators built into the vertices.
- Characterisation of the thermal expansion, including an improved determination of long-term cavity drift and the ratio of the coefficients observed in different directions.

3.2 DESIGN UPGRADES FOR REQUIRED ENVIRONMENTAL TESTING

The design upgrades planned in this project are focused two important goals: gaining confidence in achieving a rigid mounting arrangement that does not allow for cube misalignment during vibration and shock testing and developing a mounting design that enables the realisation of acceleration sensitivities matching or exceeding those recorded with the original cubic cavity. Two paths to developing such a design have previously been identified and will be investigated further within the project.

The first is using kinematic mounting principles to fully constrain the cube with ball-vee contacts at three of the vertices, with the fourth contact providing a restoring force. Such an arrangement constrains the cube to a single orientation, thereby removing the potential to misalign during shake and shock testing. This offers the benefit of being able to lightly hold the cube, which in turn means that it will see lower variation in contact force. However, the additional machining adds complexity to the manufacture and, importantly, breaks the symmetry of the design, which could have consequences for the acceleration insensitivity of the end system.

Modelling of the original cube geometry showed first order insensitivities to inertia and suggested theoretical minimum sensitivities of $\sim 10^{-17} /g$ [15]. This low sensitivity is due to the symmetrical nature of the cube and its mounting scheme. A further insensitivity to holding force was then demonstrated through truncation of the cube vertices to a specific depth that is found with FEA and known to be dependent on cube size and mirror geometry. A layout for fully constraining the cube with three ball-vee contacts was presented in [16], in which it was also demonstrated that the design can also be made insensitive to squeeze through the truncation of the vertices.

However, models assessing the sensitivity to inertia - which would be expected to increase relative to the original design - are yet to be developed. This will be completed in this project and will allow for an analysis of the acceleration sensitivity increase due to the break in symmetry caused by such a design. This will lead to an important decision point: if the design is found to induce acceleration sensitivities greater than those already achieved by the existing design then it is clear that the concept is not feasible and so efforts will be focused on the existing design. If the design is found to be feasible, further assessment will investigate

sensitivity to alignment and machining tolerances. It is also acknowledged that the machining of the vee-grooves may raise local stress, and assessment of this situation will be undertaken and design changes implemented if necessary.

A second approach will consider the positional tolerance on cube placement during the assembly. While inertial sensitivities of $2 \times 10^{-12} /g$ are predicted considering uncertainties in the material properties and manufacturing and assembly processes, experimental results have at times been higher by a few times this value. Previous projects have considered these uncertainties and have found that the most significant contribution comes from the accuracy to which the cube is centred within its frame [1]. However, only an analytical assessment of misalignment in one direction has so far been completed. This project will seek to better explore the sensitivity to cavity positioning through modelling and provide fuller understanding of the effect of that misalignment in all axes has on both sensitivity to inertia and squeeze.

While misalignment in one axis is known to introduce inertial sensitivity in the same direction, it is not currently understood how the misalignment in all axes influences the overall inertial sensitivity of the cube. Characterising the parameter space (simultaneous misalignment in axial and two orthogonal directions) will provide a better understanding of the performance of existing systems. The CCU is an interesting case, as a dual axis cubic cavity acceleration sensitivity of two axes can be measured simultaneously, meaning experimentally measured values may help to identify other factors that need consideration, such as material anisotropy. This analysis will consider different cavity spacer dimensions, as well as the effect of symmetry-reducing mirror and compensation ring combinations. Depending on these findings, additional damping or feed-forward acceleration compensation techniques between the cubic cavity chamber and platform may also be recommended.

It is then important to understand that the cube is fully constrained against displacement during environmental testing. Existing analytical models will be revisited and improved upon, including with 2-D axisymmetric models of the contact to determine an accurate equivalent stiffness for the non-linear spring supports. A full contact FEA model incorporating the cube and support spheres will be used to determine the relationship between holding force and loss of contact due to inertia, which will be used to calculate minimum allowable contact forces for future assemblies. Acceleration causing cube rotation will also be considered with this model, again building on earlier analytical calculations. Finally, the level of accelerations experienced by the cube due to different shock levels, which demonstrate that the cube support structure can maintain contact with the cavity during such shock events, will be investigated.

As part of this activity, we will also detail the required shake and shock environmental test plans that will be required in a future build.

3.3 IMPROVED CHARACTERISATION OF THE SPACER THERMAL EXPANSION

Initial results from the CCU showed that the cubic cavity ULE material is not isotropic – this is not unexpected but the level of anisotropy relating to our application has not been measured previously. So, for example, the ratio of the drift correction in orthogonal directions was not unity but measured to be ~0.7. The constancy of this ratio will need to be determined over time as part of this project. We have already demonstrated that using the clock laser to determine the ULE drift, and then compensating for this drift for the other lasers works well in the short term. However, this will only remain effective over several years if this anisotropy remains constant over this same period and may be a function of the number of times the cavity chamber is let up to atmospheric pressure.

A further effect resulting from ULE anisotropy is that the thermal expansion coefficient is zero at slightly different temperatures in the two directions. We have observed a ~2°C difference

between these two directions; for MW-CCU we anticipate controlling the ULE temperature close to the value where the effective CTE is zero for the clock axis. This would mean operating at around 2°C from the zero CTE temperature for the other axis, which will result in a non-zero linear coefficient of thermal expansion for that axis. However, with a re-determination of the ratio of thermal expansion between the axes, we do not anticipate this to be a significant problem. Even if the ULE were perfectly isotropic, we note that small changes in the temperature distribution within the ULE would ultimately limit the effect of correcting for the drift in one direction by monitoring the drift along an orthogonal axis. However, within this project, it would be useful to have more accurate measurements of the zero thermal expansion temperatures in the two cavity directions. This could be done via a static determination (i.e. via a temperature change followed by at least a day of re-thermalisation before a new frequency measurement).

The combination of ULE anisotropy and difficulties experienced at PTB in establishing temperature control of the cavity indicate that further thermal insulation will be required, given that the satellite platform is unlikely to have a high degree of temperature control. This will most likely also be required to meet typical clock laser frequency stability requirements (e.g. parts in 10^{16} over timescales from 0.1 s to 100 s). Furthermore, it would be advantageous to increase the heater capacity to allow control over a wider external temperature range. Suppression of thermally driven frequency excursion by a measure and feed-forward scheme may also be considered.

In this project, we will also consider ways of improving the effect of the thermal compensation rings described in [17]. To achieve optimum frequency stability of the clock laser over timescales of up to 100 s, it is of significant advantage to be able to operate the CCU at a temperature where the effective CTE of the clock laser cavity axis is zero. However, it is difficult to achieve this in practice if this temperature is below room temperature (e.g. < 20°C). However, [17] shows that the use of larger mirrors should improve the effect of the rings. We note that 2" diameter fused silica super-polished mirrors are available from Coastline Optics and we will consider the trade-off between larger mirrors and consider the effect of these mirrors and annuli on vibration insensitivity. We will also investigate with CMS (ThorLabs) the feasibility of contacting crystalline coatings onto this larger substrate.

3.4 REDUCING THE SWAP FOOTPRINT

A key part of the requirement for SWAP reduction is to review the option for a smaller (e.g. 3-cm) cubic spacer. Within this project, we will consider the effect of opting for a smaller cube on:

- The expected increase in frequency instability at 1397 nm as a result of a degraded thermal noise figure for the shorter cavity along the clock axis. We will consider the knock-on effect on the strontium lattice clock achievable frequency stability.
- The effectiveness of the thermal compensation rings for the situation where we continue to use 1" (25.4 mm) mirrors with 3-cm cube. In this situation, we can expect that the rings should compensate better than for a 5-cm cube for the change in effective CTE when using fused silica mirrors and a ULE spacer.
- For this smaller cube and 1" diameter mirrors and thermal compensation rings, we will model the required vertex cut-out to achieve minimal vibration insensitivity.
- The electronics design (section 0), because higher frequency EOM drive frequencies will be required to bridge the frequency interval of half the cavity FSR (i.e. half of ~5 GHz, or 2.5 GHz).

4 INTEGRATING THE CLOCK STABILISATION SYSTEM INTO THE CCU

Since the commercially available lasers that we are planning to use have a larger linewidth [14], we envisage using a hybrid electronics arrangement that is principally digital but incorporates a high bandwidth analogue servo. The feed-forward arrangement will also change, in that previously the clock axis cavity drift was measured by a slow digital lock of the laser to cavity offset frequency, where the clock laser was fixed in frequency by reference to an ultra-stable laser which serves as the NPL optical master oscillator. In the new configuration where the clock laser is stabilised directly to the cavity without prior stabilisation, the cavity drift information must be collected by another means, which will involve the counting of an RF beat signal; either with an optical frequency comb or to another reference laser. The control software must be changed to accommodate this new data source. In operation with an atomic reference the information is obtained by the logged offset required to maintain the laser at the atomic clock centre frequency. Additional methods to de-drift each cavity axis through a measure of the differential drift will also be explored.

4.1 ELECTRONICS UPGRADES

As discussed above, the main design change to the electronics will be driven by the requirement to use the CCU to stabilise the frequency of the clock laser, necessitating a higher bandwidth servo. A fully digital but lower bandwidth (measured to be in closed-loop ~30 kHz) PDH servo is being developed by NPL within another ESA project. However, these electronics which are being developed for HSL2 (contract 4000116220; see [18, 19]) do not require frequency tuning to a strontium transition frequency which we implement using the scheme in [7]. To demonstrate the feasibility of a combined analogue/ digital arrangement, we will build and test an analogue output PID stage with 1 MHz-plus servo bandwidth within this project.

Additional to the servo upgrade at 1397 nm, redesigns will be proposed that focus on simplifying the electronics and making larger frequency offsets more straightforward to generate. This upgrade aligns with the potential reduction in cube size to 3 cm. The currently used Analog Devices AD9910 DDS chips will be replaced by PLL-based wideband synthesiser chips such as the Analog Devices ADF4356. In doing so we will likely give up the electronic sideband locking method in place of a dual sideband modulation scheme in which, for simplicity, the required offset and PDH modulation is applied as two independent RF sources. Where appropriate, we will import digital servo technology from HSL2; (ESA contract 4000116220; see figure 6) into the CCU design. The HSL2 specification, however, does not include the need to tune the lasers to a specific output frequency, since only the frequency stability is specified.

5 FINITE ELEMENT ANALYSIS OF CCU CAVITY DESIGN & MOUNTING

The document details the finite element analyses conducted to verify the design choices for the Multi-Wavelength Clock Control Unit (MW-CCU) developed within ESA contract 4000132316 [20]. The CCU comprises a dual axis version of the NPL-patented force-insensitive cubic cavity rigidly mounted within an aluminium frame and a vacuum chamber. These analyses centre around verification of the holding arrangements used to maintain vertex support of the cubic cavity under expected loading conditions. Attention is also given to the concept of a vee-groove mounting arrangement for the cube by specifically considering the effect that breaking the symmetry of the cube has on the sensitivity to inertia. Further, considerations are also undertaken in respect of possible reduction of the cubic cavity size.

Firstly, the existing tetrahedral mounting arrangement is revisited to clarify the underlying cause of variation seen in the acceleration sensitivity of assembled cavity systems. Calculations and models used to justify the compressive holding force are shown, and the relative merits of alternative mounting methods are investigated. This includes discussion of the analytical and numerical modelling approaches for the tetrahedral Hertzian contact mounting configuration using PEEK spheres, and presentation of the expected values of contact force required to maintain the cubic cavity alignment under load. The material considerations of Hertzian contacts are discussed before an assessment of alternative mounting configurations is presented. Secondly a cavity size trade-off study is presented, illustrating the impact on the cube vertex truncation depth and the effective cube coefficient of thermal expansion (CTE) of reducing the cube size to 4 cm or 3 cm.

5.1 ORIGINAL CUBE GEOMETRY

To understand the variation in acceleration sensitivity achieved since the first cubic cavity system was conceived and assembled, it is instructive to revisit the original finite element modelling of the system and check consistency of results. The most detailed analysis of the concept is presented in [15], since which the design has evolved to include a variety of mirror geometries, thermal expansion compensation annuli and the use of dual bores as cavity axes. Recreation of initial results provides a means of validating current modelling practices, such as meshing, and allows consideration of some of the limitations of the original modelling in relation to estimation of sensitivity to inertia.

Within [15], the effect of numerical modelling errors was estimated through a mesh convergence study in which the cubic cavity system was modelled with and without symmetry in both COMSOL and Abaqus packages. In these models displacement at the centre of each mirror was calculated while an inwards compressive force of 1 Newton was applied at each of the mounting vertices. Where symmetry planes dissect the mounting points then the force applied is corrected accordingly. Running this simulation repeatedly while sweeping through the cube corner truncation depths enables determination of the null cut depth, the depth at which the cube becomes insensitive to the mounting force on the vertices, from plots such as Figure 17 and those presented in [15].

While a correctly truncated cube will be insensitive to squeeze, some error in truncation depth must be assumed. This is handled by comparing the sensitivity to a 1 Newton mounting force at an assumed cut depth, for example 6.5 mm in Figure 18, across models of varying mesh density to find a range within which the sensitivity is expected to lie. The mesh densities in the model affect its number of degrees of freedom, defined as the number of nodes multiplied by the number of variables to be solved. Such data also provide a validation of the model mesh, with a clear convergence seen in the higher density meshes. To demonstrate that current simulations produce results consistent with original efforts, this data, originally presented in [15], has been supplemented with results from more recent models and is presented in Figure 18. Model parameters match those presented in [15] and are repeated in **Table 2** for convenience.

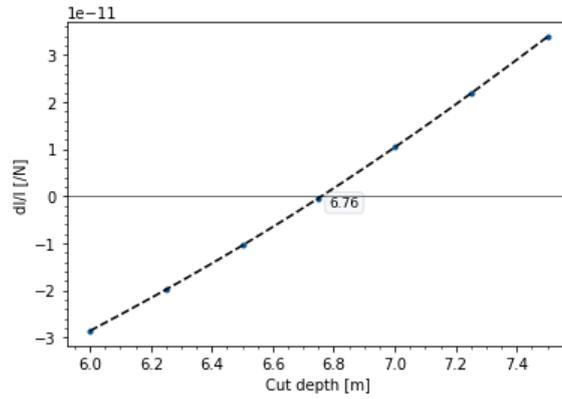


Figure 17 - Response of the cavity to squeeze as a function of the cut depth at the cube vertices

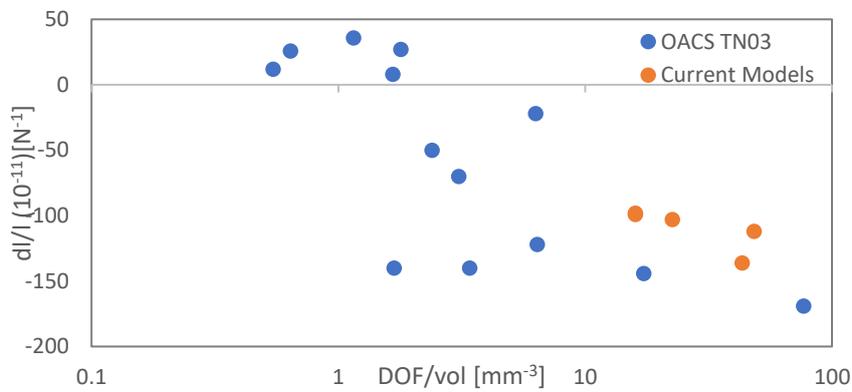


Figure 18 - Length change for a cut depth of 6.5 mm as a function of the number of model degrees of freedom per unit volume. Original data is taken from [15]

Table 2 - Geometric and material properties used in the cut depth calculation finite element model

Geometric Parameter	Value [mm]	Material Parameter	Value
Cube size	50	Young's Modulus (ULE)	67.6 [GPa]
Vertex truncation depth	6.5	Poisson's Ratio (ULE)	0.17
Cube bore diameter	5.1	Density (ULE)	2210 [kg m ⁻³]
Mirror radius	6.34	Young's Modulus (FS)	72.7 [GPa]
Mirror thickness	4.02	Poisson's Ratio (FS)	0.16
Mirror radius of curvature	500	Density (FS)	2201 [kg m ⁻³]
Support radius	0.28		

5.2 SENSITIVITY TO INERTIA

Sensitivity to inertia, as experimentally characterised through ‘inversion tests’ in which frequency shifts are monitored while the cavity is inverted about each of its axes, is a critical performance characteristic of the cubic cavity. The design has been demonstrated to achieve extremely low passive acceleration sensitivities; however, consistently achieving the level of $2.5 (\times 10^{-11} \text{ g}^{-1})$ originally demonstrated has proven challenging (see Table 3). This section is dedicated to better understanding the cause.

In particular, models are developed to predict the effect mounting force has on the cavity's sensitivity to axial inertia. This is primarily in response to the observation that recent cavity systems have been assembled with much higher holding forces than the original (best performing) system. For example, a normal force of ~650 N was applied to each of the vertices when assembling the OSRC cavity, compared to just 100 N for the original system.

Table 3 - Acceleration sensitivities of each of the feeder project cavities (data from [1] [21] [22])

Project	Acceleration sensitivity ($\times 10^{-11} \text{ g}^{-1}$)			Holding force (N)
	Axial	Transverse	Transverse	
[1]	2.5	0.2	0.01	<100
HSL2	10	8.6	35	Bonded ULE
OSRC	8	10	0.2	~650
CCU clock (aux.)	2.74 (0.24)	1.71 (3.00)	10.00 (9.56)	~1500

As above, models predicting sensitivity to inertial force are described in most detail within [15]. In these models a body load of ρg is applied to the sub-domains of the model while the supports, modelled by a circular patch of a defined radius, are fixed. Once again, the displacement at the mirror centres is output from the model to calculated length change of the cavity for the given inertial load.

Defining the supports in this way within the models provides a convenient method of constraining the cavity; however, it presents several issues. Considering the above question on the effect of holding force on inertial sensitivity, it is clear that by fixing the support patches it is not possible to study this because any force applied within the FEA models is against a fixed boundary and so will have no effect. Therefore, an alternative set of modelled boundary conditions for the cube must be found if we are to demonstrate the effect that holding force has on inertial sensitivity. A second concern regarding the use of fixed constraints at the contact points is that it is not truly physically representative of the real system in which the cube is free to detach from its supports. This departure from reality can be identified when the modelled cube deformations are made visible by applying a scaling factor, as illustrated in Figure 19. Here we see an unexpected stretching of the cube at the contact point, which although small is large relative to cavity length changes.

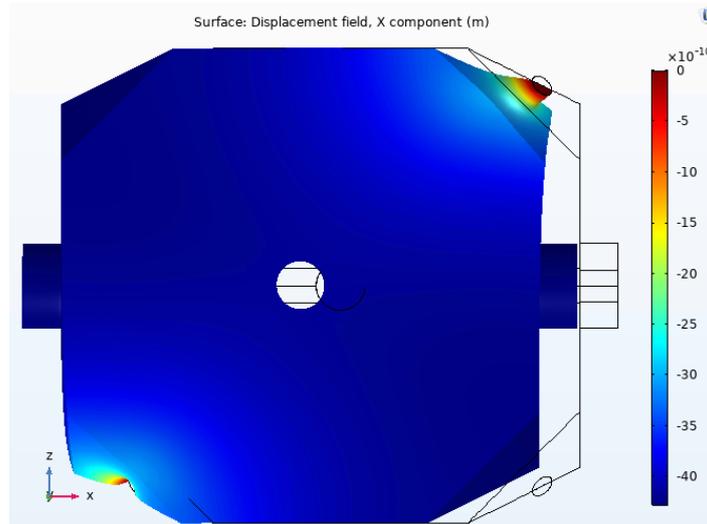


Figure 19 – Deformation of the cavity under axial inertia. The supports are fixed and a body load due to gravity (acceleration of g) is applied in the $-x$ direction. The deformed cube geometry is given in colour, while a wireframe represents the geometry without loading. The deformation is scaled by a factor of 1,000,000 to make it visible. The cube is seen to stretch at the mounting point due to the fixed boundary condition.

An alternative approach is to apply constraints known as Spring Foundations at the mounting boundaries. Spring Foundations connect the structural boundary to a fixed ground with a stiffness that is defined by the user. Opting for a spring type of total force as a function of extension and applying a pre-deformation to the spring allows us to closely replicate the behaviour of the tetrahedral contacts without resorting to contact modelling.

An equivalent stiffness for the contact sphere mountings can be found from Hertzian contact theory as discussed in [15] and [23]. The extent of compression, Z , is known as the distance of approach and is defined as the displacement of the sphere minus its compression:

$$Z = \frac{a^2}{r} = \left(\frac{4}{3} \frac{k}{E^*}\right)^{\frac{2}{3}} P^{\frac{2}{3}} r^{-\frac{1}{3}}. \quad (1)$$

Here, a is the radius of the contact circle, r is the radius of the contact sphere, E^* relates to the Young's modulus and Poisson's ratio of the cube and sphere materials, P is the force with which the sphere is pressed into the plane, and k is a term defined as

$$k = \frac{9}{16} \left[(1 - \nu_1^2) + (1 - \nu_2^2) \frac{E_1}{E_2} \right] \quad (2)$$

where E_2 is the Young's modulus of the indenter (in this case the PEEK sphere) and ν_1 and ν_2 are the Poisson's ratio of the specimen and indenter respectively. Rewriting the expression for Z in the form of a spring force equation,

$$P = k_t Z^{\frac{3}{2}} \quad (3)$$

relates the force at the contact to the distance by which the sphere has been compressed, as shown in Figure 27 later in the Section. These equivalent stiffness parameters, k_t , are introduced into the model above and summarised in **Table 4**.

Table 4 – Equivalent stiffness parameters

Symbol	Equation	Value	Description
k_t	$1e10 [N/m^{(3/2)}]$	$1E10 N/m^{(3/2)}$	tangential spring constant
u_t	$0 [m]$	$0 m$	tangential spring predeformation
k_n	$6.38e8 [N/m]$	$6.38E8 N/m$	spring constant of peek sphere
u_n	$-101.28e-6 [m]$	$-1.0128E-4 m$	displacement of sphere
f_n	$k_n u_n^{(3/2)}$	$-650.29 N$	force of Hertz contact

Introducing a Spring Foundation constraint in this way links the force at the contact patch to displacement of the cube without fixing the cube geometry. A body load due to gravity can then be applied to observe the displacement of the mirrors due to external accelerations. The resulting deformation is illustrated in Figure 20, in which the contact patch is now seen to compress with the normal force of the support while the whole geometry is translated due to a body load of 1 .

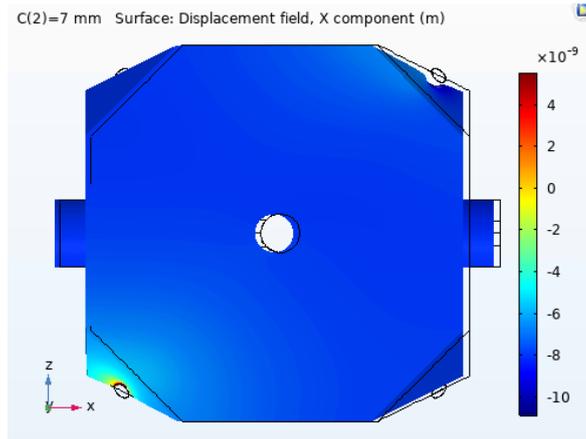


Figure 20 – Deformation of the cavity under axial inertia. The supports are constrained with Spring Foundations and a body load due to gravity (acceleration of g) is applied in the $-x$ direction. The deformed cube geometry is given in colour, while a wireframe represents the geometry without loading. The deformation is scaled by a factor of 100,000 to make it visible.

To identify the effect that squeeze has on the cube sensitivity to inertial forces, the relative displacement between the centre of the cavity mirrors is output from the model and divided by the cavity length to find a fractional length change, dl/l . This is plotted against cut depth for a range of mounting forces and presented in Figure 21. The result, which includes the presence of a $1 g$ body load in the $-x$ direction, is familiar as the trend seen when determining the cut depth of an unloaded cube geometry.

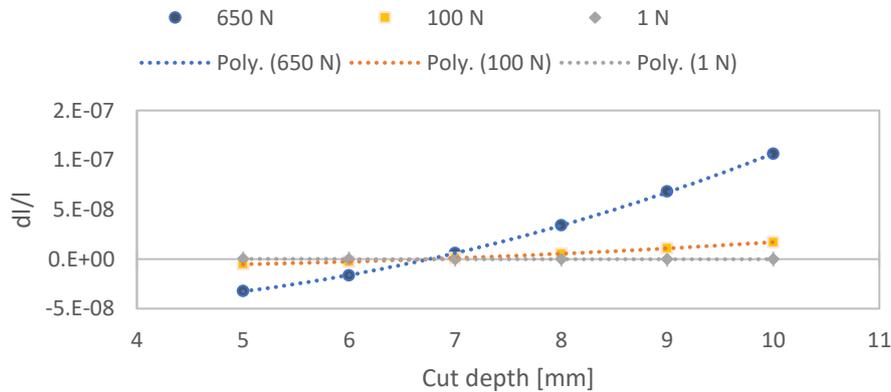


Figure 21 – Cavity fractional length change as a function of cut depth in the presence of a $1 g$ body load.

This data alone does not identify the inertial sensitivity of the cavity. To determine the change in cavity length due to the body load on the cavity it is important to subtract the fractional length change found when no body load is applied, i.e. the cavity length change due to the mounting forces. Here it is critical to note that this deformation is non-zero except at the precise null cut depth; this is in contrast to the previous method employed in which the contact patch was fixed (i.e. not loaded). Calculation of the fractional length change of the cavity without the presence of a body load show near identical results to those presented in Figure 21, regardless of the

body force. For the case of a normal mounting force of 650 N at each of the vertices, the length change due to the introduction of a 1 g body load is shown in Figure 22.

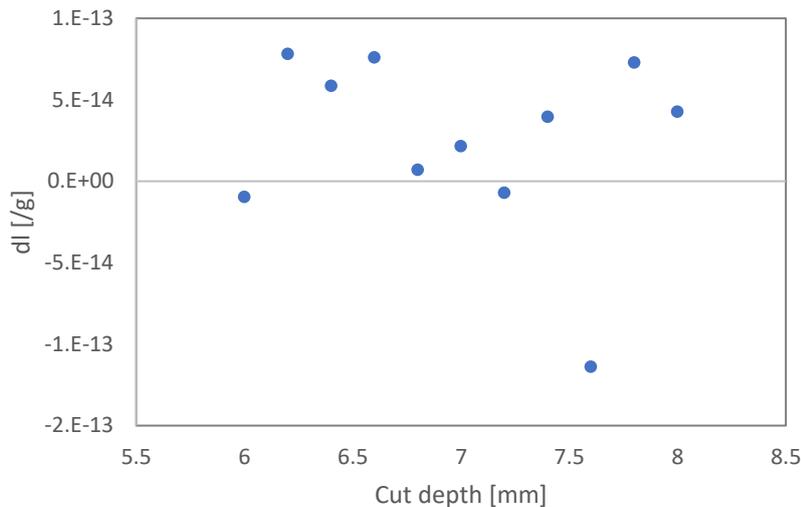


Figure 22 – Cavity length change due to 1 g body load as a function of cut depth. The cube is held with a normal force of 650 N.

The residual length change does not exhibit any largescale trend, supporting the suggestion in [15] that axial inertia is negligible regardless of the corner truncation depth. The scatter in the results may suggest further mesh refinement would be beneficial; however, even with an estimated sensitivity equal to this range the contribution to acceleration sensitivity would be at the 1.92×10^{-13} /g level, almost two orders of magnitude lower than the most insensitive cavity axial performance measured. Therefore, models do not support that an increased holding force would negatively affect the acceleration sensitivity of the cubic cavity system as measured with inversion tests, regardless of the cut depth tolerance.

5.3 MINIMUM HOLDING FORCE

This subsection details the models used to assess what holding force is required to prevent a loss of contact between the cubic cavity and its mounting spheres under the worst-case loads expected during launch. Initial cubic cavity designs for ESA calculated holding forces based on a 2D approximation of the cube and assumptions around the maximum translational and rotational accelerations that the cube would experience. In terms of translational accelerations it was assumed that, because pairs of mounting spheres oppose each other, the maximum change in force at any one support would be half of the maximum acceleration load [24]. For the initial OSRC project estimates of a maximum acceleration of 45 g and cube mass, m , of 0.275 kg, this equated to a lower limit on the static pre-load required to prevent loss of contact at the mounts of 62 N.

In terms of rotational accelerations it was assumed that equal and opposite 45 g accelerations on either side of the cube were equivalent to the worst-case rotational loading (*i.e.* an angular acceleration, α , of approximately 12,000 radians / s²), as is shown in Figure 23. This angular acceleration is very large, and the validity of this assumption remains uncertain, but it represents a conservative estimate of the worst-case angular acceleration the cube mounting supports may experience.

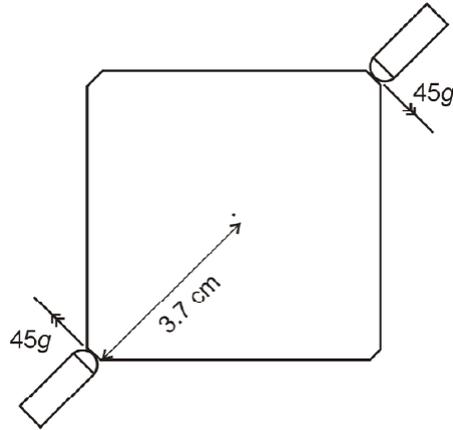


Figure 23 - Illustration of the rotational acceleration assumed to act as two equal and opposite translational accelerations, reproduced from Ref. [24]

The gravitational torque must be matched by the frictional torque provided by the supports for the cavity not to slip *i.e.*

$$\tau_g = I\alpha = \frac{mL^2}{6} \alpha \quad (4)$$

where τ_g is the gravitational torque, I is the moment of inertia and L is the cube edge length. For the 5 cm cubic cavity under these assumptions this correlates to a required frictional torque of at least 1.4 Nm. A friction coefficient, μ , of 0.1 (nylon on glass) was assumed, and for four spherical mounts this means that the total frictional force

$$F = 0.4P \quad (5)$$

where P is the applied pre-load at each of the four supports. As the individual components of the frictional force each act at 3.7 cm from the cube centre, this equates to a total frictional torque

$$\tau_F = 0.015P \quad (6)$$

which means at least 93 N is required to prevent slipping of the cube. Making a conservative assumption by adding this to the force required to prevent loss of contact meant each support needed to be pre-loaded with a compressive force of at least 155 N.

Using these analytical calculations finite element models could be developed to assess the stresses that arise in the cube when under compression. Finite element models are constructed using COMSOL Multiphysics v5.5. To model the Hertzian contacts at the supports small circular patches were defined on the supported vertices (see Figure 24), with radius a defined from Hertzian theory [23] as

$$a^3 = \frac{4}{3} \left(\frac{k}{E_1} \right) Pr \quad (7)$$

where here E_1 is the Young's modulus for the specimen (in this case the ULE glass cube), P is the applied compression load, r is the radius of the indenter and the k term is defined earlier. Spring foundations were applied to these patches to model the mounting support, with the spring constants related to the predicted required compressive forces.

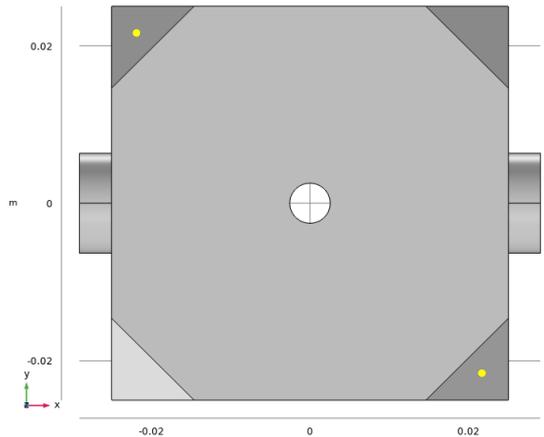


Figure 24- Top-down view of the full 3D model showing two small contact patches (highlighted in yellow) at opposite vertices where spring foundations have been applied to model the Hertzian contacts

5.4 CONTACT MODELLING

Using equivalent stiffness models gives a quick approximation of the support performance but this approach has several key limitations. Firstly, frictional effects and kinematic contacting between the spheres and cavity cannot be modelled. This means that the amount of resistance to rotational acceleration provided by the supports cannot be verified accurately, and uncertainties remain over the accuracy of the translational acceleration sensitivities. Additionally, no mechanism exists for modelling any loss of contact, and a spring force would be erroneously recorded at a spring foundation patch where the cube is being pulled away from the support due to acceleration loads. Furthermore, the contact patch area cannot change under loading as it is defined from the applied contact force, meaning any change in contact pressure, and thus sphere compression, is not captured.

Contact models of the cubic cavity and its mounts are constructed to ensure these effects are captured. Contact models also enable the accuracy of Hertzian theory to be assessed by comparing the predicted contact area sizes and the amount of compression, Z , that occurs at the supports when loaded.

5.4.1 2D ASSYMETRIC MODELS

The first set of models are 2D axisymmetric models with two PEEK spheres being pressed into and then released from the surface of a ULE half cube (see Figure 25). By pressing the sphere into the cube, the contact area is determined by assessing the stresses and resulting deformation of the PEEK sphere and ULE cube.

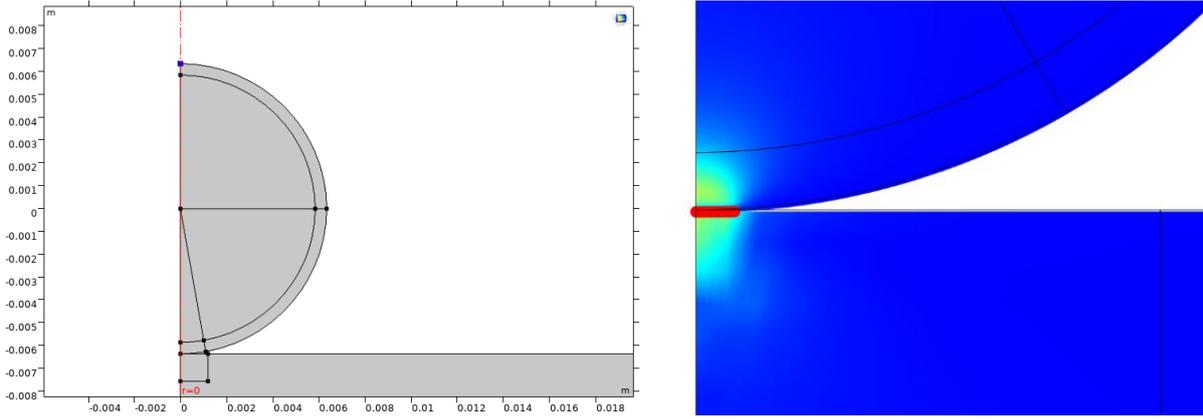


Figure 25 - Illustration of the 2D axisymmetric modelling approach. Geometry partitions are used to add finer meshes to the contact regions (left), and then the actual contact patch size (red line) and stress states can be measured during compression (right)

Figure 26 shows the how the calculated contact patch radius using these 2D models compares to Hertzian theory predictions for a range of compressive loads modelled as either a full body load, a boundary load on the outer hemisphere surface or a point load at the centre of the outer hemisphere surface. It can be seen that the Hertzian theory matches quite well to the finite element models, with the FEA approach predicting marginally larger contact areas at high compressive loads.

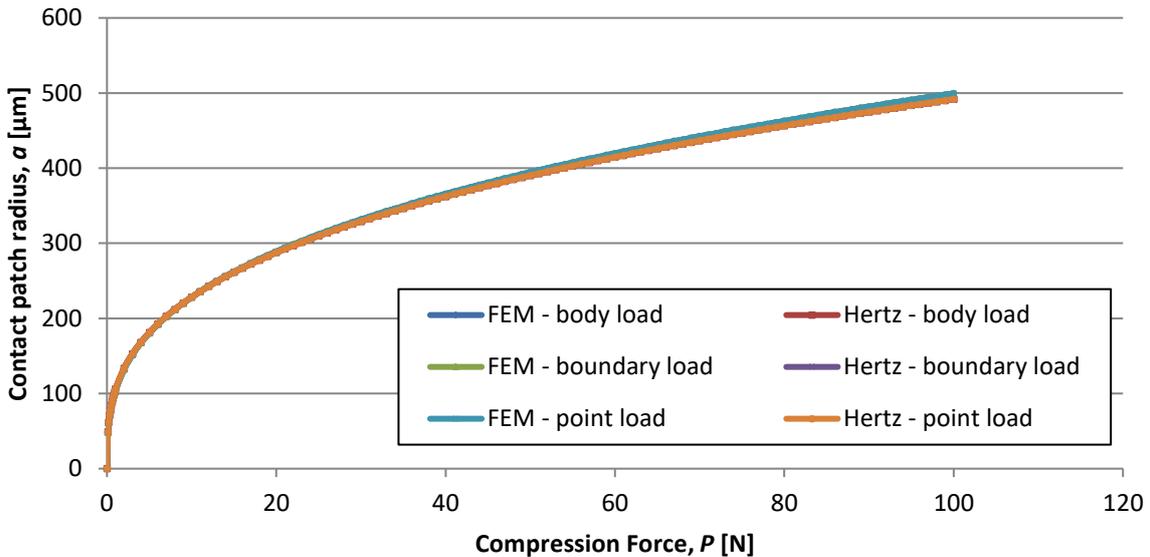


Figure 26 - Comparisons of the estimated contact area between finite element models and Hertzian contact theory

Agreement between theory and numerical solutions is weaker for the distance of approach. The distance of approach in a Hertzian contact is directly linked to the compressive force applied to the sphere. Figure 27 shows that Hertzian contact theory predicts a much higher distance of approach than the finite element model suggesting that it underestimates the

sphere compression. This mismatch between FEA and theoretical distance of approach values supports similar findings reported in the literature [25].

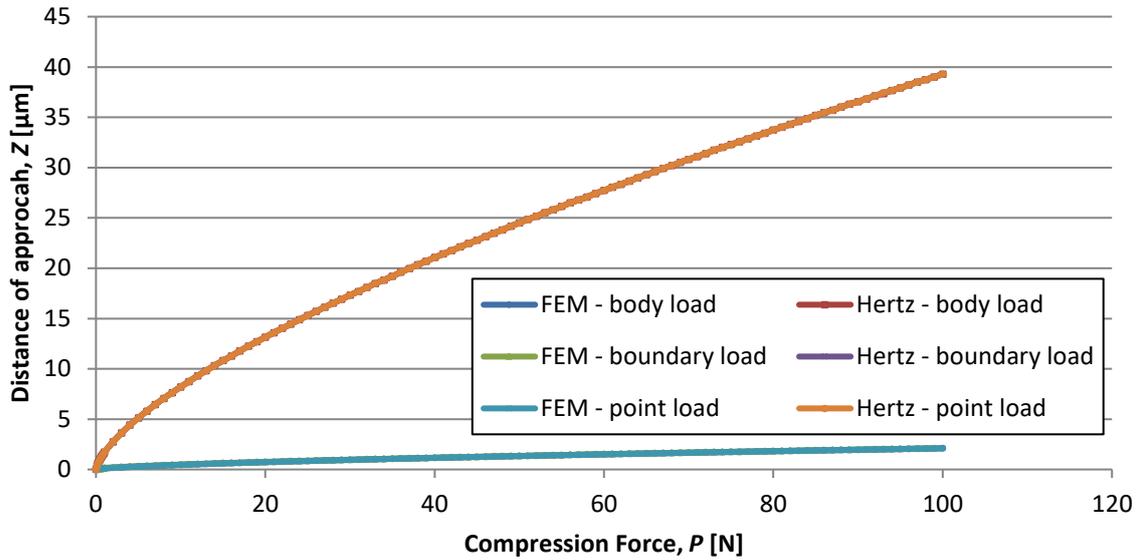


Figure 27 - Comparisons of the estimated distance of approach between finite element models and Hertzian contact theory

5.4.2 3D MODELS; CUBES 7 HEMISPHERES

The second set of models are 3D models that include the full cube geometry and four hemispheres to model the sphere and stud supports (see Figure 28). Here the hemispheres are displaced along the cube vertex normal by an amount equal to the distance of approach. For reference the geometric and material properties used in the FEA models are listed in **Table 5**.

Table 5 - Geometric and material properties used in the finite element models

Geometric Parameter	Value [mm]	Material Parameter	Value
Cube size	50	Young's Modulus (ULE)	67.6 [GPa]
Vertex truncation depth	6	Poisson's Ratio (ULE)	0.17
Cube bore diameter	5.1	Density (ULE)	2214 [kg m ⁻³]
Mirror diameter	12.68	Young's Modulus (PEEK)	3.6 [GPa]
Mirror thickness	4.02	Poisson's Ratio (PEEK)	0.39
Mirror radius of curvature	500	Density (PEEK)	1260 [kg m ⁻³]
Mounting sphere radius	6.35		

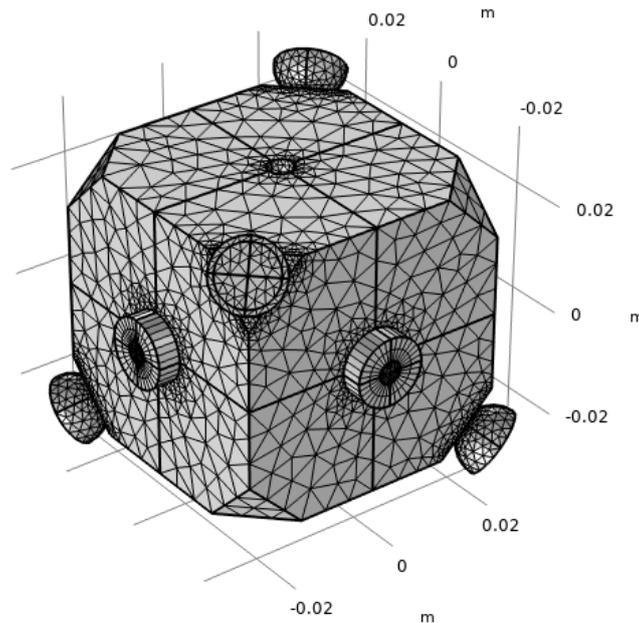


Figure 28 - View of the meshed 3D contact model used in these analyses. The mesh density is increased on the cube vertices and on the hemisphere contacting patches

For convenience the distance of approach is calculated using Hertzian theory for an arbitrary range of applied compression forces. The actual contact forces between the cube and the hemispheres are calculated by integrating the contact stresses over the cube vertices after the hemispheres have been displaced into the cube (displacements are equal to the desired Hertzian distance of approach). This displacement-controlled approach gives accurate values for the initial applied contact loads that do not rely on Hertzian theory.

While the contact patch size predicted by Hertzian theory gives good estimates, a circular patch area of double the radius ($2a$) is created on the spheres so that the contact is still captured in the event of small rotations of the cube during loaded. This patch is meshed with a much higher density than the surrounding geometry to enable accurate calculation of the contact stresses. While it is not necessary to use a contact patch on the spheres, meshing the entire hemisphere surface with a dense enough mesh would be computationally expensive. As a compromise it is set at twice the size predicted so that any small slippage or non-symmetrical deformation of the spheres (that would cause a movement of the contact patch across the vertices) is captured when the cube is subjected to g-loads. An increased mesh density is also applied to the cube vertices, although this is not as dense as at the contact patch.

A contact pair is defined between the cube vertices and the contact patches so that the two surfaces cannot pass through each other. Additionally a frictional behaviour is assigned to this, assuming the same friction coefficient as used with the nylon-glass calculations in previous projects ($\mu = 0.1$) [24]. The type of stress field resulting from this contact behaviour when the spheres are pressed into the cube can be seen in Figure 29.

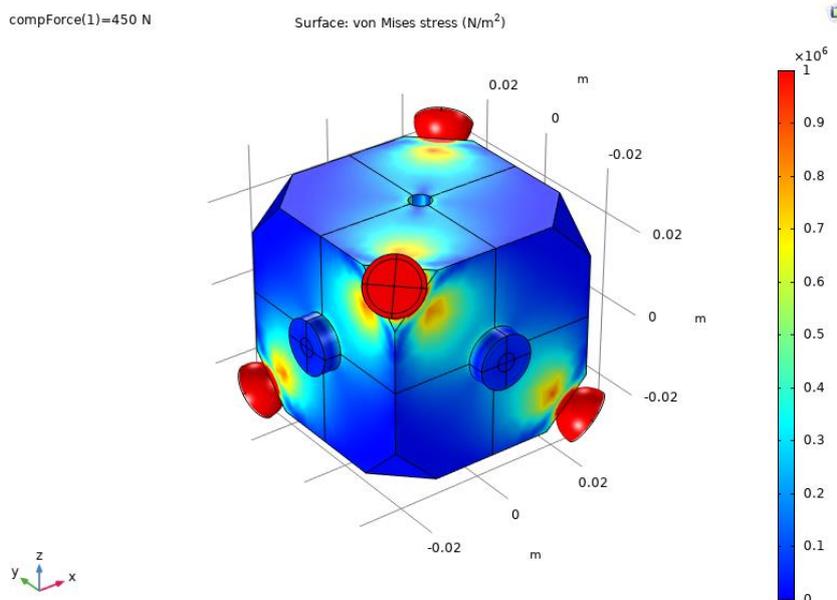


Figure 29 - Illustration of 3D model showing the stress field symmetry when hemispheres are contacted to the cavity under no gravitational acceleration

5.5 ACCELERATION MODELLING

The full 3D models are used to model the cube response to expected quasi-static launch accelerations. In previous ESA project proposals the maximum quasi-static acceleration was assumed to be 45 g [24], and previous ESA project requirements have stated quasi-static loads of 30 g in any direction, or alternatively $KQ(48 / mass^{0.25})$ g [26], where KQ is a qualification factor (equal to 1.4) defined in ECSS-E-ST-32-10C [27]. Additionally, expected levels of random vibration from previous projects had values of approximately 10 and 15 g_{rms} for in-plane and out-of-plane excitation respectively. No quasi-static requirement was set for MW-CCU [20], so using previous requirements as a benchmark (*i.e.* a quasi-static load with a margin of safety between 2-3, or a +6-sigma random vibration) a maximum acceleration load of 90 g is modelled.

Modelling the contacts is a highly nonlinear problem, and when bodily accelerations are included model convergence can become an issue. To improve the robustness of convergence an auxiliary sweep of the applied g-load is included *i.e.* the models linearly ramp the g-load in 30 g increments, using the previous ramp increment solution as the initial condition of the next ramp increment. Mesh refinements were conducted on the models in this section to ensure that the residual contact forces (or measured displacements depending on model in question) were converged to within 98% of the values reported by the finest meshes tested.

5.5.1 LOSS OF CONTACT DUE TO TRANSLATIONAL ACCELERATION

In terms of the cube losing contact with one of the hemispheres (and hence risking misalignment of the optical axis) the worst-case translational acceleration would not be parallel to one of the XYZ axes, but parallel to the cube diagonal, *i.e.* directly away from a spherical contact in a direction normal to the contact vertex. This loading condition is shown in Figure 30.

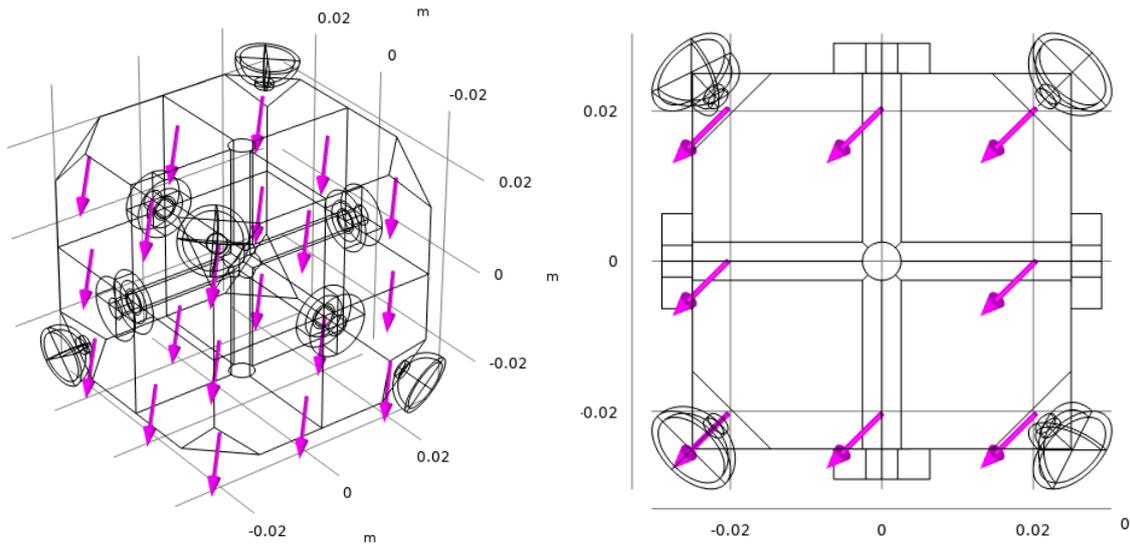


Figure 30 – Orthographic (right) and planar (left) views of the 3D cavity model showing a translational static acceleration applied diagonally. This is the assumed worst-case translational acceleration orientation as it directly pulls the cube away from a sphere contact in a direction parallel to the vertex normal.

By varying the displacement of the hemispheres into the cube and integrating the contact pressures at each hemisphere as the cube is loaded the minimum required applied support pre-load can be identified as shown in Figure 31.

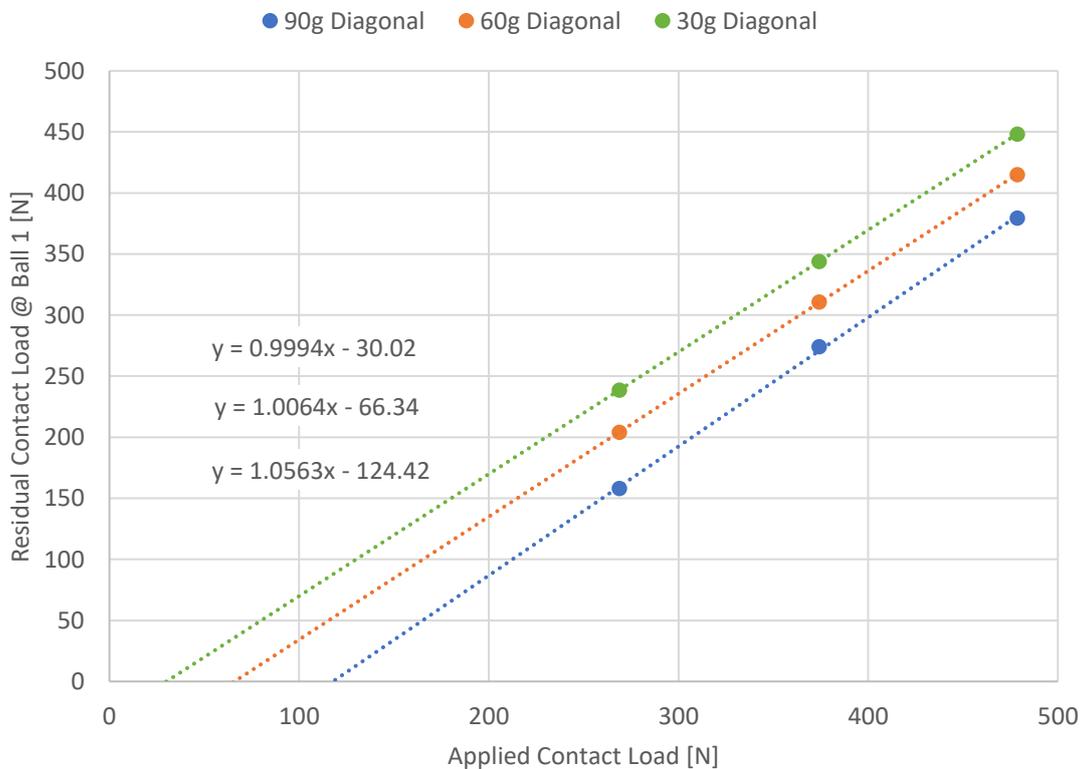


Figure 31 - Residual contact force vs. applied contact force for a range of static translational accelerations (friction coefficient $\mu = 0.1$)

The relationship between applied pre-load and residual contact load under translational accelerations was found to be linear, with the resulting trendline indicating that a pre-load of at least 30 N, 66 N and 118 N must be applied to each mounting sphere to ensure no loss of contact under translational accelerations of 30 g, 60 g and 90 g respectively.

Figure 31 shows the residual contact force trends for a case where friction was modelled at the contacts with an assumed friction coefficient of $\mu = 0.1$. Figure 32 shows the same graph but considers the impact of friction on the minimum required holding force. Given the uncertainty surrounding the actual friction coefficient for a Hertzian contact between a PEEK sphere and a ULE cube it plots the results of friction-free contact models and models including a higher friction ($\mu = 0.3$). It can be seen that the x-axis intercept – *i.e.* the minimum required contact load – varies significantly with friction, and the difference between this load for frictionless and high friction models increases with increasing acceleration magnitude.

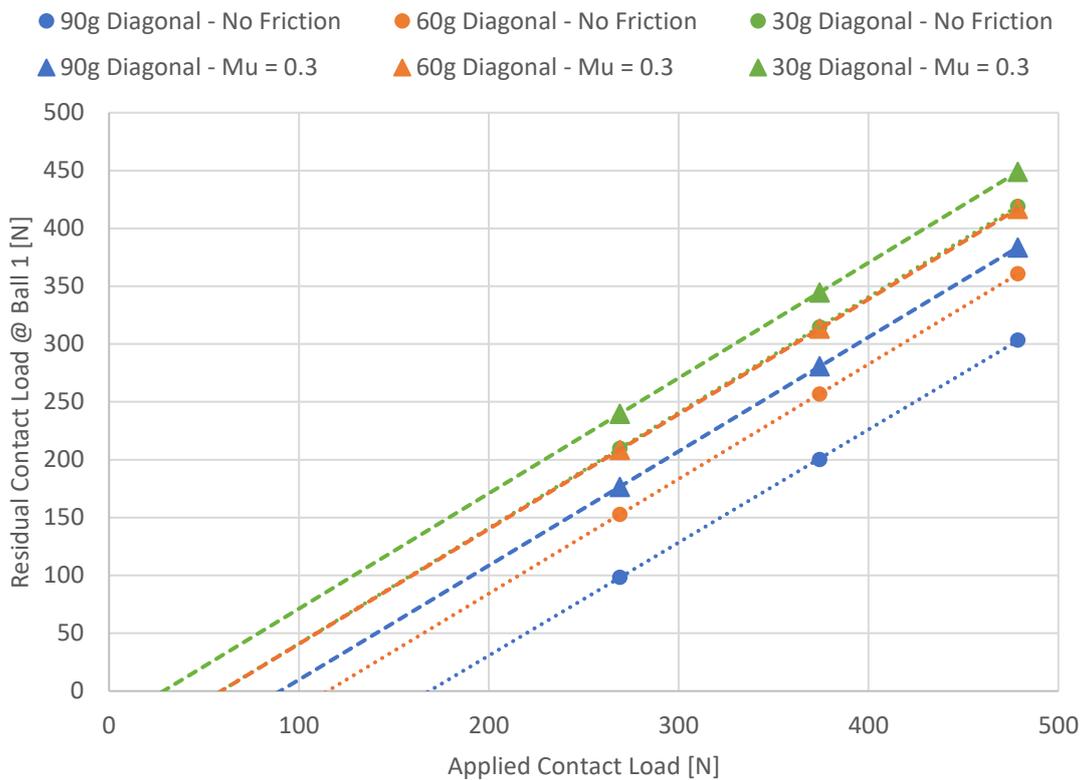


Figure 32 - Residual contact force vs. applied contact force for a range of static translational accelerations and contact friction coefficients

Table 6 shows the full range of calculated minimum holding forces needed to prevent loss of contact under this translational acceleration load case. It shows that there is a large drop-off in the required holding force when friction is included, but that increasing μ beyond 0.2 does not yield much more of a reduction.

Friction clearly plays a significant role in resisting contact loss under loading – it is not just determined by kinematic constraint. As the expected friction coefficient between the ULE and PEEK spheres is uncertain a series of characterisation experiments should be carried out to determine it, but based on this modelling it can be seen that a friction coefficient of 0.2 could perhaps be set as a target value to ensure that a minimal holding force is used at the cavity supports.

Table 6 - Minimum holding forces required to prevent loss of Hertzian contact under translational acceleration for different levels of friction. The y and x terms in the trendline refer to the axes of Figure 31 and Figure 32, i.e. residual contact load and applied contact load.

Static Translational Acceleration	Friction Coefficient	Linear Trendline	Minimum Required Holding Force [N]
30 g	0	$y = 0.9983x - 58.73$	59
30 g	0.1	$y = 0.9994x - 30.02$	30
30 g	0.2	$y = 0.9976x - 28.804$	29
30 g	0.3	$y = 0.9972x - 28.515$	29
60 g	0	$y = 0.9916x - 114.09$	115
60 g	0.1	$y = 1.0064x - 66.34$	66
60 g	0.2	$y = 0.9952x - 59.825$	60
60 g	0.3	$y = 0.9933x - 58.571$	59
90 g	0	$y = 0.9772x - 164.79$	169
90 g	0.1	$y = 1.0563x - 124.42$	118
90 g	0.2	$y = 0.9919x - 91.876$	93
90 g	0.3	$y = 0.9871x - 88.914$	90

5.5.2 ANGULAR DISPLACEMENT DUE TO TORQUE

Angular displacements of the cube are modelled in 3D using similar approximations to those described earlier, i.e. the worst rotational load case that the cube could experience is a constant torque about the non-optical principal axis (z-axis) that causes a peak tangential acceleration (90 g) at the support vertices.

Torques applied to solid bodies in COMSOL must be defined using equivalent tangential forces, F_{eq} , defined in a cylindrical coordinate system. The required torque can be expressed as

$$\tau_{Rq} = F_{eq}R = \frac{I_{zz}a_T}{R_V} \quad (8)$$

where a_T is the tangential acceleration, R is the range of radial distances (in a spherical coordinate system) at which the equivalent forces are applied, and R_V is the radial distance between the cube centre and the truncated vertices. The moment of inertia for the cube and its mirrors, I_{zz} , is calculated using the mass properties toolset in COMSOL and found to be $1.14e-4 \text{ kgm}^2$. The cube mass is similarly evaluated as 0.272 kg. For vertex tangential accelerations of 30 g, 60 g and 90 g the required constant torques are 0.902 Nm, 1.804 Nm and 2.705 Nm respectively.

Using this assumption, a conservative estimate of the pre-load required to ensure minimal rotation of the cube optical axis is calculated. This minimal rotation limit is assumed to be such that the optical axis is not permanently rotated by more than 5 arcminutes, and is based on analysis conducted by STI for the HSL2 project [28]. The rotation is calculated from the FEA solution by using the dot product to evaluate the angle between the unrotated and rotated position vectors describing the location of the mirror centres on one axis. It is assumed that the average rotation angle of each mirror describes the rotation of the optical axis.

Figure 33 shows the expected rotation of the optical axes when the cube is subjected to a constant torque and no friction is included in the contact model. This shows the resistance to rotation provided by just the kinematic restriction of the mounts, with the torque being resisted entirely by elastic deformation at the contacts. It can be seen that this alone permits rotations in the order of hundreds of arc seconds, with the largest torque causing rotations >1200 arcminutes (>20°) for the contact loads investigated.

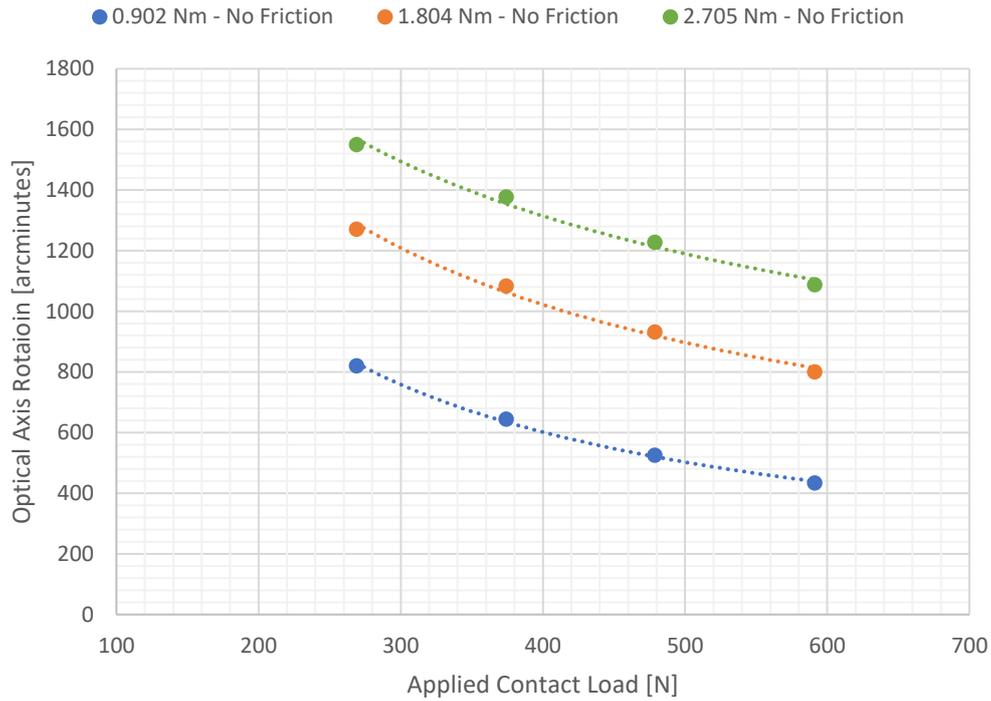


Figure 33 - Rotation of the cube optical axes due to a constant torque about the z- axis with no frictional behaviour at the contacts

Figure 34 shows the same information but for models where two different values of friction ($\mu = 0.1, 0.3$) are included. The expected rotations are significantly reduced by over two orders of magnitude, illustrating that the contribution to rotation from ‘slipping’ of the Hertzian contact is critical.

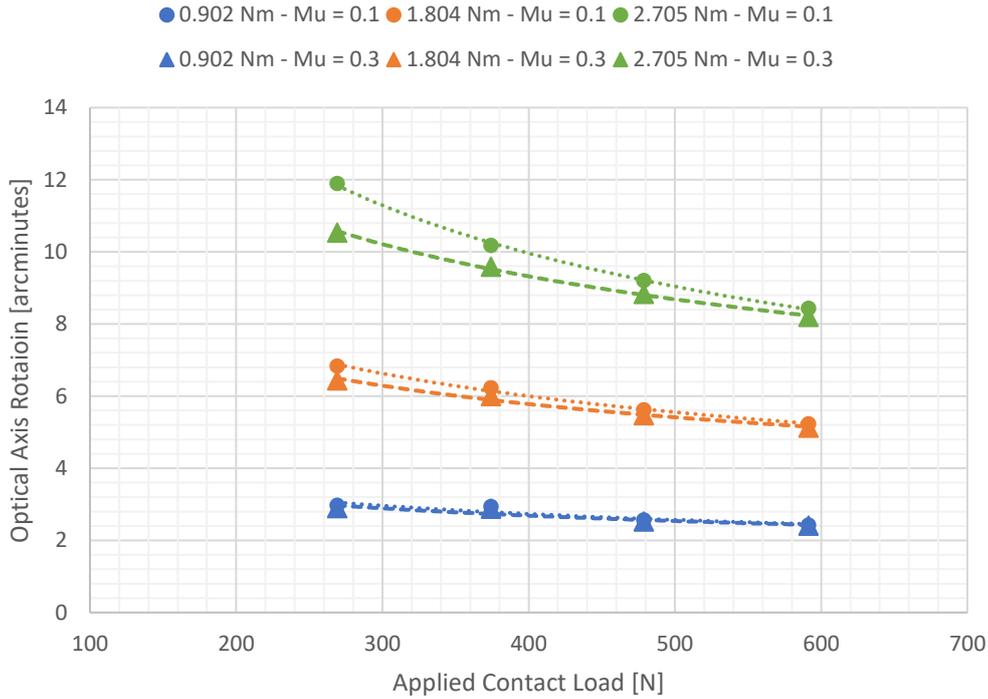


Figure 34 - Rotation of the cube optical axes due to a constant torque about the z-axis with frictional behaviour included at the contacts

The potential benefit from friction coefficient values larger than 0.1 appears to be less than for the case of diagonal translational acceleration loading. Additionally, the trend between rotation and applied contact load appears to be nonlinear (power law relations were assumed for the trendlines in Figure 33 and Figure 34). The required contact force to limit rotation to <5 arcminutes is around ~600 N for the 1.804 Nm torque and is suggested that the required contact force could be in the order of several kN for the 2.705 Nm torque, although this extrapolation relies on the power law assumption being valid at kilonewton load levels (which were beyond the scope of this current work). **Table 7** lists some expected rotations for the range of contact forces and friction values modelled.

Table 7 – Predicted optic axis rotations under 300 N and 500 N contact loads for different levels of friction. The y and x terms in the trendline refer to the axes of Figure 33 and Figure 34, i.e. optical axis rotation and applied contact load.

Constant Torque [Nm]	Vertex Tangential Acceleration	Friction Coefficient	Trendline	Rotation @ 300 N contact [arcminutes]	Rotation @ 500 N contact [arcminutes]
0.902	30 g	0	$y = 75144x^{-0.806}$	757	502
0.902	30 g	0.1	$y = 14.44x^{-0.278}$	2.96	2.57
0.902	30 g	0.2	$y = 12.678x^{-0.258}$	2.91	2.55
0.902	30 g	0.3	$y = 12.335x^{-0.254}$	2.90	2.54
1.804	60 g	0	$y = 33932x^{-0.585}$	1206	895
1.804	60 g	0.1	$y = 47.494x^{-0.345}$	6.64	5.57
1.804	60 g	0.2	$y = 35.867x^{-0.303}$	6.37	5.46
1.804	60 g	0.3	$y = 33.549x^{-0.294}$	6.27	5.40
2.705	90 g	0	$y = 19068x^{-0.446}$	1498	1193
2.705	90 g	0.1	$y = 135.14x^{-0.435}$	11.30	9.05
2.705	90 g	0.2	$y = 69.873x^{-0.334}$	10.40	8.77
2.705	90 g	0.3	$y = 62.46x^{-0.317}$	10.24	8.71

Given that the highest minimum required holding force to resist translational loss of contact was 118 N (when friction was included), the required holding forces to constrain rotations are large. It is worth noting that the torques values are conservative: it is unlikely that a purely rotational load case would occur causing such large tangential accelerations at the vertices, and it is more likely that launch accelerations of these magnitudes would occur in a translational manner.

Furthermore, the rotations measured from the FEA results are taken by applying the torques as quasi-static loads, which is itself a conservative assumption, as it is likely that some rotation is ‘recovered’ (*i.e.* not all the rotation is permanent) upon unloading of the cube, with the amount recovered depending on the applied load and frictional forces. Re-computing the analyses while including a final step that removes the applied torque gives an indication of how much of the rotation is elastic (and recoverable) and how much is non-recoverable ‘slip’ under a 90 g load.

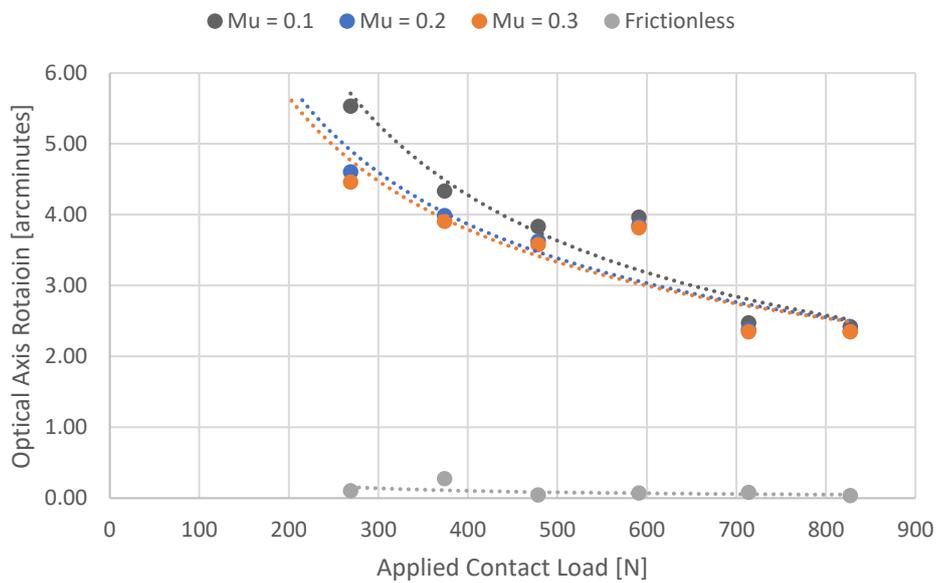


Figure 35 – Residual rotation of the cube optical axis after torque causing 90 g vertex tangential acceleration is removed

While friction reduces the expected rotation when the cube is loaded, it also causes a proportion of this rotation to become non-reversible after unloading, although it can be seen that this is typically in the order of several arcminutes of rotation. Figure 35 shows the predicted residual optical axis rotation once the maximum torque (2.705 Nm) has been removed from the cube, representing the aftermath of one maximum rotational load event. These FEA results exhibit a degree of unexpected variation, perhaps suggesting that the mesh – which is currently refined for modelling the Hertzian contacts – could be further refined for the rest of the cube (although the models are already computationally expensive). However, when assuming a power law trend with these results it can be seen that relatively low contact forces are needed to constrain the residual rotation to <5 arcminutes: a force of 323 N should enable this for the lowest value of friction expected (see **Table 8**).

Table 8 - Minimum holding forces required to constrain optic axis rotation to <5 arcminutes for different levels of friction

Friction Coefficient	Trendline	Minimum contact load required to restrict residual rotation to <5 arcminutes
0	$y = 53.479x^{-1.042}$	0
0.1	$y = 337.14x^{-0.729}$	323
0.2	$y = 141.01x^{-0.6}$	261
0.3	$y = 120.95x^{-0.578}$	248

5.5.3 MINIMUM HOLDING FORCE

In the OSRC project, it was estimated that the minimum required holding force each contact would need to provide as the sum of the forces required for sufficient translational and rotational constraint (155 N) [24]. This assumes a very conservative worst-case scenario whereby the cavity is subject to both the maximum translational and rotational loads at once. Based on the FEA results conducted in this work package, the required contact force should be 441 N for maximum accelerations of 90 g with contact friction coefficients of 0.1. This is higher than in earlier analysis simply because higher accelerations are considered so to be more conservative. If the Hertzian contact friction coefficients are found to be higher at 0.3, this total contact force reduces to 338 N. Given that the force required to resist rotation is consistently much higher than that required to resist translation it could be argued that the minimum contact force could be taken as only that required to resist rotation, giving forces of 323 N or 248 N for friction coefficients of 0.1 and 0.3 respectively.

5.5.4 EIGENFREQUENCY ANALYSES

Figure 36 shows the first and second groups of eigenmodes of the mounted cubic cavity: a group of three cube rigid-body-rotation modes and a group of three cube rigid-body-translation modes. In this case, a cube-sphere contact force of 269 N is applied. Including the rest of the cavity structural assembly would capture the stiffnesses of the sphere mounting studs, which may slightly alter the expected eigenfrequencies.

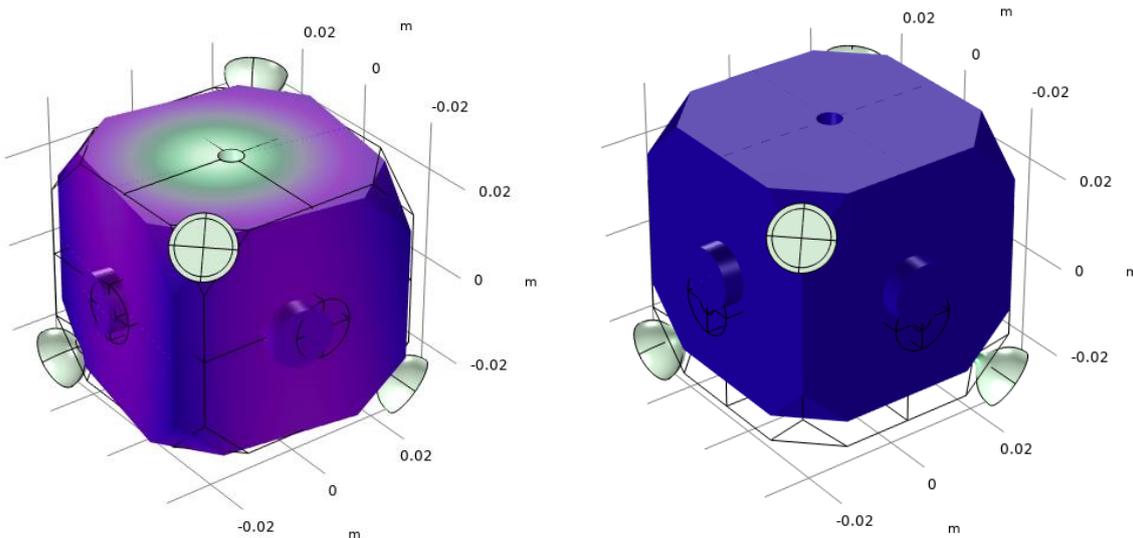


Figure 36 - Eigenfrequency analysis of the mounted cavity with a contact load of 269 N. The first three eigenmodes (cube rotation) occur predicted to occur between 80-82 Hz (left), and the second three (cube translation) between 822-823 Hz (right).

The first three rotational eigenmodes occur at relatively low frequencies, although the magnitudes of these frequencies increase with applied contact load. Figure 37 shows the FEA predictions for the lowest eigenfrequencies and assumes a power law trend in the data. Previous ESA projects have required the first eigenmode of the system to occur at no less than 140 Hz [29], and the trend in Figure 37 suggests a contact force of 792 N or higher will be required to achieve this.

It is worth noting that these low frequency eigenmodes are rotational, and that shock qualification of the cubic cavity systems is conducted by applying loads only in a translational manner. COMSOL eigenfrequency analyses calculate effective modal masses for each eigenmode, which describe how much each mode is affected by accelerations with high effective modal mass indicating that a load in that direction will have a high contribution to the eigenfrequency. Effective modal masses are listed in **Table 9** for these first six eigenmodes when the cube is held with a relatively low contact force (269 N) and a relatively high contact force (827 N). They are expressed as percentages of the system mass for translational cases, or respective moment of inertia for rotational cases ($m = 0.2716 \text{ kg}$, $I_{xx} = I_{yy} = 1.127 \times 10^{-4} \text{ kgm}^2$ and $I_{zz} = 1.143 \times 10^{-4} \text{ kgm}^2$).

It can be seen that while the rotational effective mass percentages are very high for the first three eigenmodes in each load case (and that increasing the contact force increases the first eigenfrequency), the translational percentages are negligible. This suggests that the translational accelerations used to validate the performance of the cubic cavity systems will have negligible contributions to these rotational modes which, in turn, suggests that the high required contact force suggested by Figure 37 may be overly-conservative. The cubic cavity system developed for the ESA OSRC project utilised a similar Hertzian contact mounting arrangement (with 650 N mounting forces) and, when tested, the first eigenfrequency of the entire system (including dampers, vacuum chamber and heat shields etc.) was predicted to be driven by dampers attached to the vacuum chamber [30], and was found in testing to exist at over 160 Hz [29]. This further suggests that the low rotational eigenfrequencies found in these FEA analyses will not compromise the system performance under expected vibrational and shock tests.

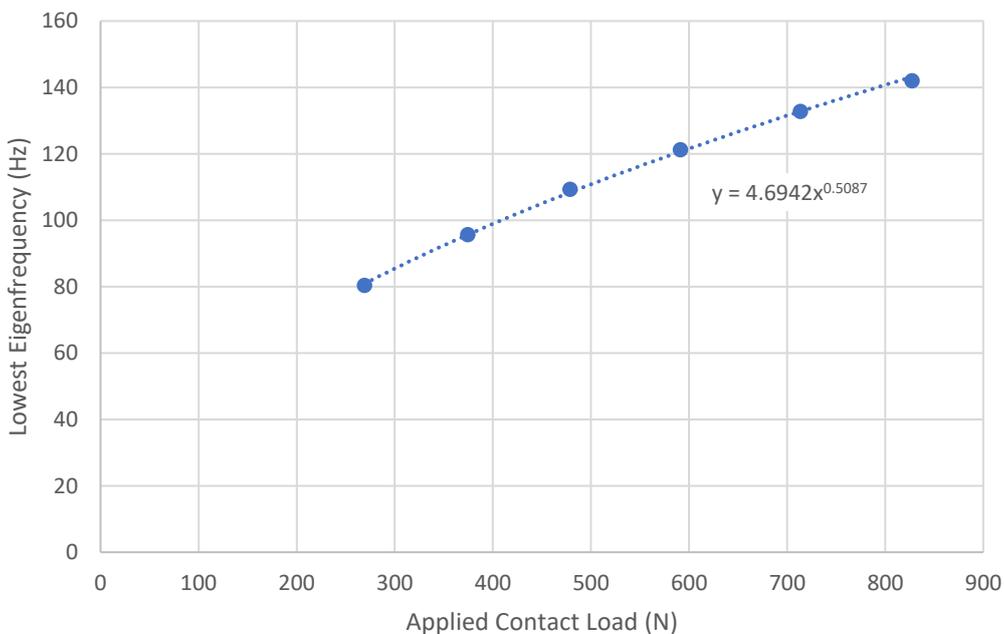


Figure 37 - Lowest eigenfrequency of the mounted cube vs. applied contact load between the cube vertices and the mounting spheres

Table 9 - Participation factors for the first six eigenmodes with cube contact forces of 269 N and 827 N

Applied Contact Load (N)	Eigenfrequency (Hz)	Effective Modal Mass (Translational) as % of Cavity Mass			Effective Modal Mass (Rotational) as % of Cavity Inertia		
		X	Y	Z	X	Y	Z
269	80.351	0.0	0.0	0.0	4.3	0.2	96
269	81.037	0.0	0.0	0.0	27	73	0.5
269	81.077	0.0	0.0	0.0	69	27	4.0
269	822.32	59	1.4	40	0.0	0.0	0.0
269	822.33	8.5	66	26	0.0	0.0	0.0
269	822.77	32	33	35	0.0	0.0	0.0
827	141.98	0.0	0.0	0.0	0.8	0.1	99
827	142.99	0.0	0.0	0.0	25	75	0.6
827	143.07	0.0	0.0	0.0	75	25	0.3
827	1000.5	52	0.6	48	0.0	0.0	0.0
827	1000.7	0.1	99	0.7	0.0	0.0	0.0
827	1000.9	48	0.2	51	0.0	0.0	0.0

5.6 CUBE FRACTURE

The material used for the cavity spherical mounts must be able to transmit the required compression force while being soft enough to avoid fracturing the cube. A loaded Hertzian contact causes a tensile stress field to occur on the vertex surface, with the level of tensile stress varying with radial distance from the contact point. Initial estimations based on Hertzian theory [23] were conducted in the OSRC [24] project, with maximum tensile stress

$$\sigma_m = \frac{1}{2}(1 - 2\nu_1)p_0 \tag{9}$$

where p_0 is the average pressure between the sphere and the cube

$$p_0 = \frac{P}{\pi a^2} \tag{10}$$

Other terms in these equations are defined previously in Section 0. Figure 38 shows the expected maximum tensile stresses calculated this way for half-inch diameter mounting spheres. The red line indicates the estimated fracture stress for ULE glass at a Hertzian contact (based on literature stress values shown in **Table 10**), and the black line shows the compressive pre-load used in the OSRC cavity (which survived environmental testing) [29].

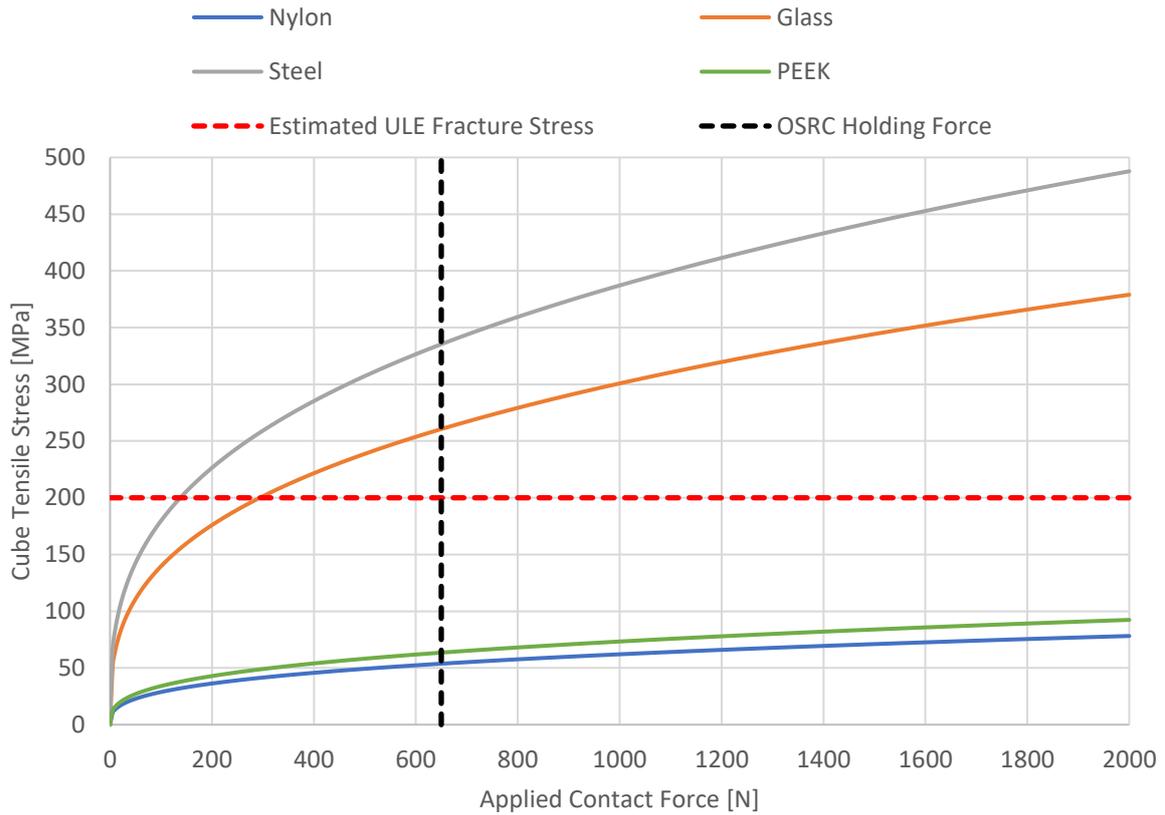


Figure 38 - Maximum tensile stresses resulting from 1/2-inch diameter spherical mounts (from [23]).

It can be seen that under the assumptions of Hertzian theory it is not expected that the cube will fracture as the PEEK spheres are too soft induce a stress close to the fracture stress for the expected compressive forces. Although nylon appears to give even better performance, PEEK is preferred for space-based cavity systems due to its superior outgassing performance.

Table 10 - Hertzian contact tensile failure stress values from literature. Adapted from Ref. [32]

Cavity material	E_1 [GPa]	ν_1	Sphere material	E_2 [GPa]	ν_2	r [mm]	P [N]	a [mm]	p_0 [GPa]	σ_m [MPa]	Ref.
Glass	70	0.25	Glass	70	0.25	2.5	105	0.17	1.1	276	[31]
Glass	70	0.25	Glass	70	0.25	5	231	0.29	0.9	226	[32]
Glass	70	0.25	Steel	210	0.28	6.35	500	0.35	1.3	326	[23]

5.6.1 Contact Size Effects

Literature has consistently suggested that the tensile strength of glass under a Hertzian contact load is significantly higher than its bulk tensile strength (usually reported around 50 MPa for ULE glass [33]). Due to the uncertainty surrounding this strength value it is worth noting that the expected tensile stress can be further reduced by using mounting spheres with larger diameters. Figure 39 shows the predicted tensile stresses using spheres with 1-inch diameters indicating that even at the large compressive load of the OSRC project the expected tensile stress can be below even the bulk ULE tensile strength.

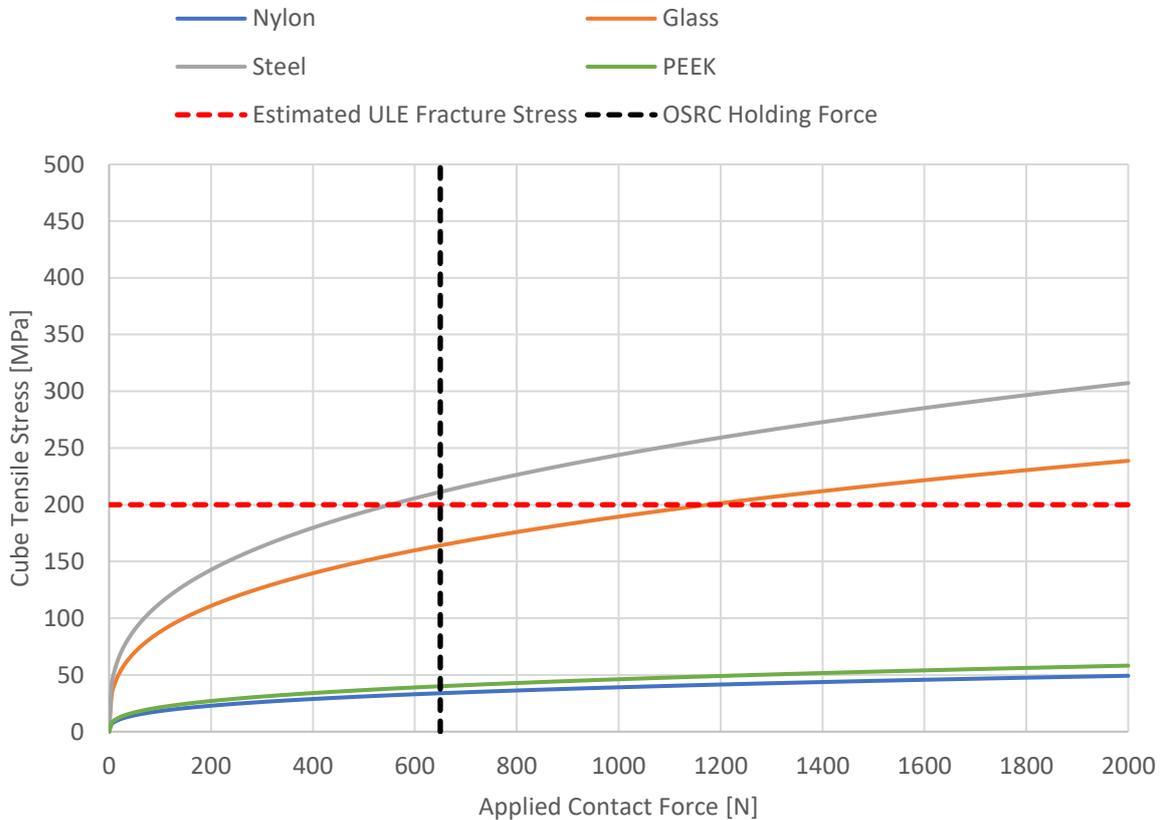


Figure 39 - Maximum tensile stresses resulting from 1-inch diameter spherical mounts. (from [23]).

5.7 PATHWAYS TO REDUCING INERTIA SENSITIVITY

5.7.1 Ball-Vee Contacts

Several concepts for alternative mounting arrangements were initially discussed in [34]. Concepts were proposed as possible geometries that might require a lower holding force without increasing the risk of losing alignment of the cube during shake testing, with the overall goal being to find a solution that demonstrates acceleration insensitivity at or below the level originally reported.

Of the proposed concepts a single geometry has been developed for further analysis: the use of vee-grooves to locate the cube. This concept follows the principle of exact location to constrain the cube with six points of contact (traditionally a cone-vee-flat arrangement, but also commonly achieved with three vee-ball contacts). With the Hertzian contact nested within vee-grooves and the cube exactly constrained (i.e. having only a single orientation in which all six contacts are engaged) such a cube is not expected to suffer from a risk of displacement during shake testing and in the case of loss of contact of any points would return to the original position due to the full constraint.

Initial modelling considered whether a null sensitivity to squeeze could be attained for the geometry; however, a full model of the sensitivity to inertia was not considered. Here the initial model is developed to include assessment of sensitivity to axial inertia – the cube sensitivity that is experimentally tested through inversion testing. The meshed geometry is presented in Figure 40.

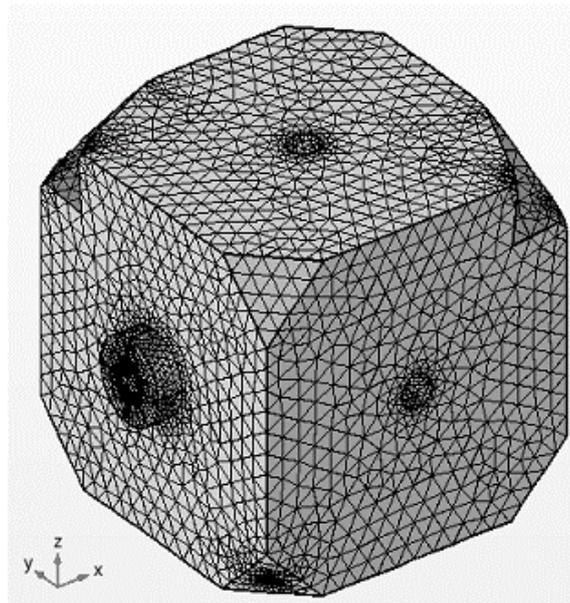


Figure 40 – Illustration of meshed single axis cubic cavity with ball-vee contacts at three of the four tetrahedral supports giving six points of contacts (and a retaining contact) to exactly constrain the cube and prevent slippage. This comes at the cost of broken symmetry.

In the model, three vee-groove contacts are implemented as three of the tetrahedral contacts, while the fourth remains flat and acts as a retaining force holding the cube into the six points of contact. An acceleration force of $1g$ is applied in the X-direction while all contact points are constrained as fixed. From the calculated mirror centre displacements, a cavity length change of $2.75 \times 10^{-12} m/g$ is found, giving a fractional cavity length change of $5.14 \times 10^{-11} /g$. This result suggests that the addition of vee-grooves into the support structure comes at a high cost to the inertial insensitivity of the cavity, and the results posted by the original truncated cube geometry could not be matched with such a design.

It is worth noting that further investigation could be possible here, in particular studies could consider the effect of various ball sizes and angles for the vee-groove. However, fundamentally the reduced symmetry of the design would be expected to increase the sensitivity to inertia and so other solutions to minimising acceleration sensitivity should first be explored.

5.7.2 Cube positioning

Here we consider whether the holding force is responsible for an increase in acceleration sensitivity seen in recently assembled cubic cavity systems. Models incorporating holding force into characterisation of inertia sensitivity suggest that increased holding forces do not increase the systems inertial sensitivity, while Table 3 shows that there is a lack of correlation between holding force and acceleration sensitivity in measured systems. Previous analyses of the uncertainties have suggested that the position of the cube relative to the frame it is held in are the most significant factor in its acceleration sensitivity performance, but these are not substantiated with numerical modelling. Here, a model to characterise the effect of positional error on the cube sensitivity to axial and transverse inertia is developed and illustrated in Figure 41

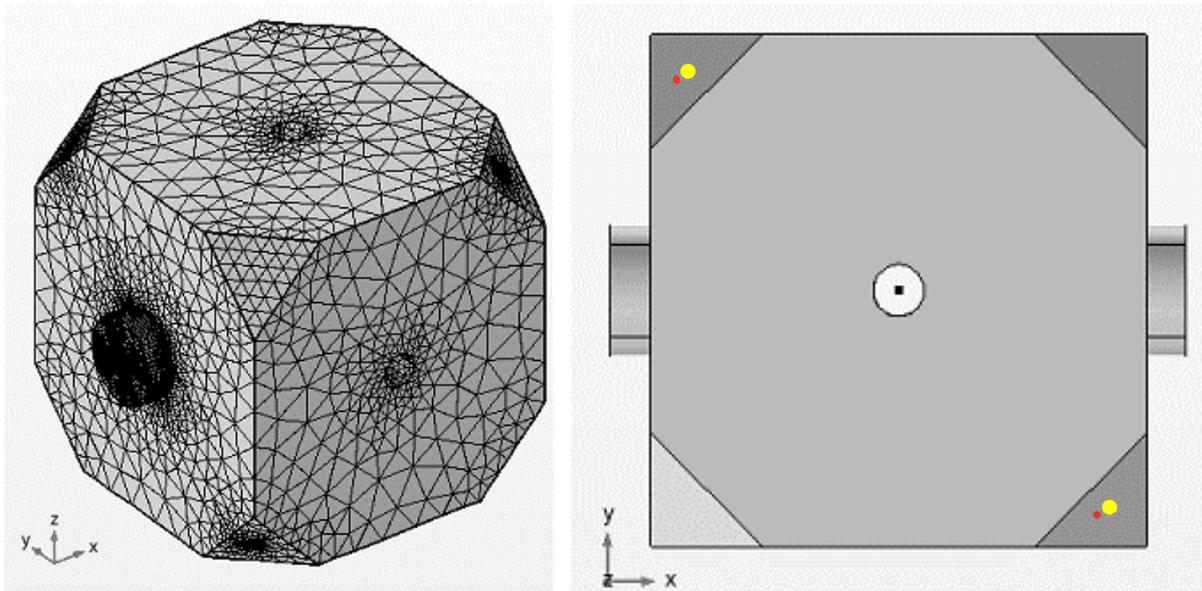


Figure 41 – Illustration of a truncated single axis cubic cavity with Hertzian contact positions dictated by errors in the cube location relative to the centre of its frame. Left: meshed geometry. Right: illustration of the Hertzian contact points (red dots) displaced from the centre of the vertices (yellow dots).

In this model the location of the Hertzian contact on the truncated vertex is displaced to account for positional misalignments in the cube position along the X, Y and Z axes. The case of a body load applied in the X and Y axes are considered to determine the effect of axial and transverse inertia respectively, and the results summarised in **Table 11**. Results indicate that misalignment outside of the axis of the body load has little effect on the cube sensitivity, while misalignment along the axis of the body load is largest of any contributor to inertial sensitivity (as detailed in [15]).

Table 11 - Inertial sensitivities of the cube for positional errors of 500 μm in the X, Y and Z axes

errX (mm)	errY (mm)	errZ (mm)	Gravity vector	dI/L (/g)	Inertia
0	0	0	X	-1.96E-13	Axial
0	0	0.5	X	-1.82E-14	Axial
0	0.5	0	X	-4.11E-13	Axial
0	0.5	0.5	X	-1.17E-12	Axial
0.5	0	0	X	-7.51E-11	Axial
0.5	0	0.5	X	-7.46E-11	Axial
0.5	0.5	0	X	-7.52E-11	Axial
0.5	0.5	0.5	X	-7.62E-11	Axial
0	0	0	Y	2.17E-13	Transverse
0	0	0.5	Y	-4.2E-13	Transverse
0	0.5	0	Y	4.39E-11	Transverse
0	0.5	0.5	Y	4.37E-11	Transverse
0.5	0	0	Y	-4.8E-13	Transverse
0.5	0	0.5	Y	1.73E-12	Transverse
0.5	0.5	0	Y	4.38E-11	Transverse
0.5	0.5	0.5	Y	4.5E-11	Transverse

These results are in good agreement with the analytical solution [35], $\Delta L/L = \rho g \epsilon / 2E$, where ρ and E are ULE density and Young's modulus and ϵ the positional error. Experimental effort is ongoing to estimate the positional accuracy which is achievable with the current assembly procedure and implement improvements to this process. Given the inability of numerical models to find causes to explain the increase in acceleration sensitivity, characterising the accuracy to which the cubes can be assembled is considered a critical next step.

6 CUBE SIZE TRADE-OFF

6.1 CUT DEPTH

Another aspect of the current modelling work is to establish the trade-offs associated with smaller cube geometries. A focus here is to understand the effectiveness of thermal compensation annuli as a function of the relative size between the annuli and the cube. These annuli act to correct the shift in the zero coefficient of thermal expansion of the cube, typically selected to be close to room temperature, due to the use of mirrors of dissimilar material such as fused silica. The use of fused silica for the mirror substrate material is preferred over ULE due to a lower material mechanical loss factor which results in a lower cavity thermal noise floor [36]. To investigate this trade-off a $\frac{1}{4}$ symmetry model employing of the 5 cm cubic cavity with 1" mirrors and annuli has been built, representing a dual-axis configuration.

6.2 MESH REFINEMENTS CHECKS

A mesh refinement study akin to that presented in Figure 18 was conducted to provide confidence in the cut depth values being obtained, and the geometry was then modified to 3 cm and 4 cm cube cases. Determining whether there is a suitable geometry to give a null to squeeze force is an essential first step after which the thermal behaviour can be modelled, as the use of thermal annuli (assumed to be 5 mm thick in this study) is only relevant if the cube is first proven to be insensitive to squeeze.

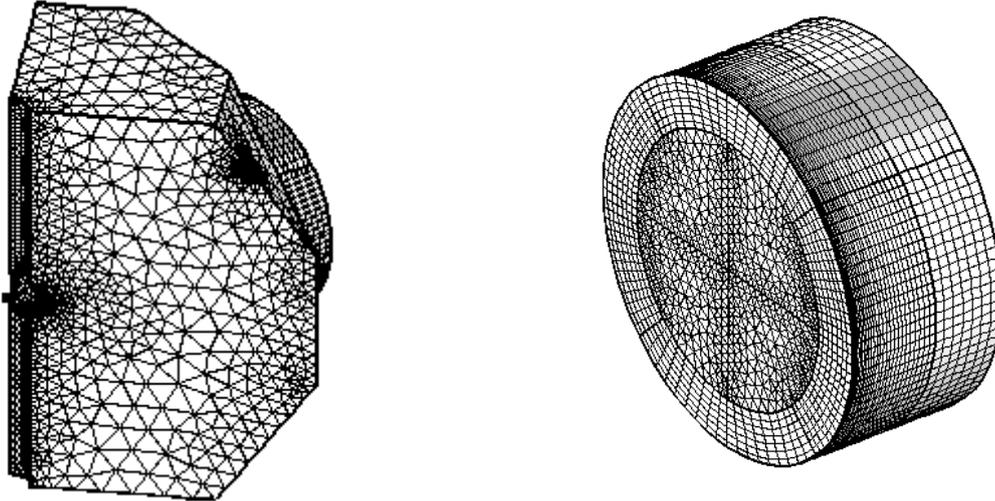


Figure 42 – Illustration of meshed dual axis cubic cavity geometry. Left: $\frac{1}{4}$ geometry with symmetry planes $y = x$ and $y = -x$ with mesh refinement at the truncated vertices. Right: Mirror and geometry with mesh refinement.

From earlier mesh refinement studies, it has been found that the most influential region of the mesh is that of the mirror, particularly near the mirror centre where the cavity deformation is considered. To illustrate the significance of the definition of the mirror mesh relative to the bulk cube mesh the two have been meshed with independent resolution, providing cut depth results for combinations of coarse and fine meshes on the two domains. The results of this study are presented in Figure 43 in which the resulting cut depth is seen to converge with increasing mesh density (increased degrees of freedom per unit area). In the figure the varying mesh resolution of the bulk cube is indicated by the colour of the points. As the mesh density increases (e.g. from 'normal' to 'fine') it is found that only a minor change to the calculated cube response is found. In contrast, increasing the density of the mirror mesh has a significant effect on the calculated cube length change. This can be understood when considering the four points with the highest mesh density, all of which are closely converged regardless of the mesh refinement on the cube. The results are also summarised in **Table 12**.

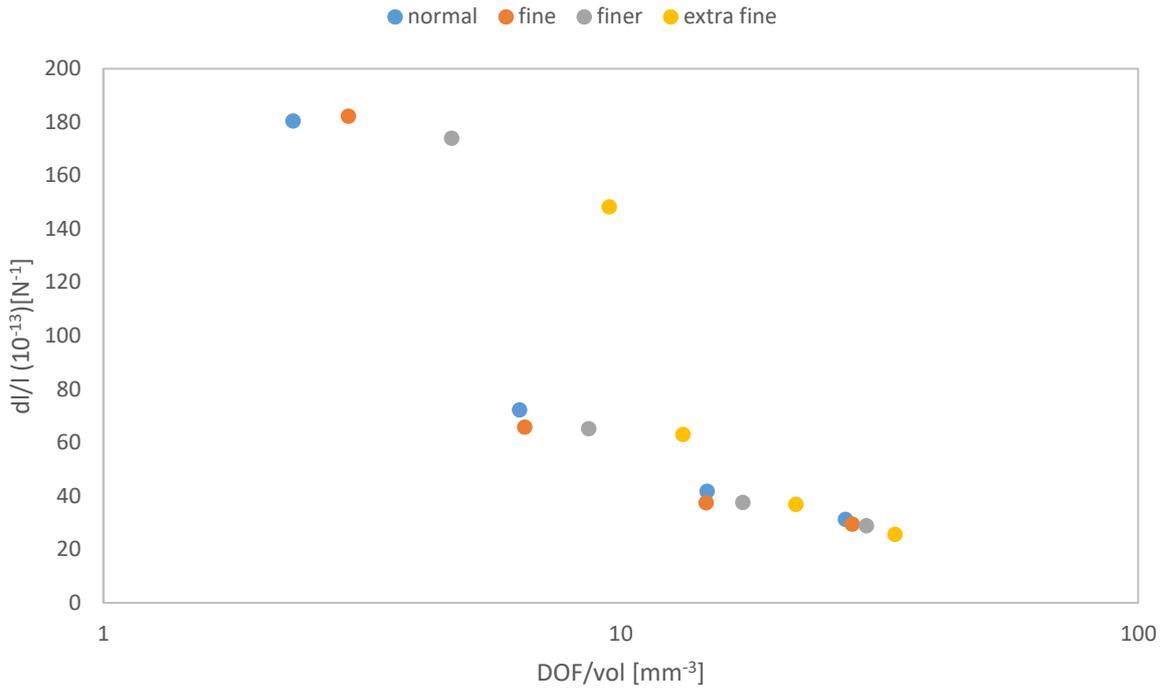


Figure 43 – Finite element modelling results of mesh convergence study of a dual axis cubic cavity geometry with 1” mirrors and annuli. ¼ symmetry is employed to reduce the mesh size required (see geometry illustrated in Figure 42).

Table 12 – Dual axis cubic cavity mesh refinement study parameters (data also presented in Figure 43). Here the bulk mesh densities ‘normal’, ‘fine’, ‘finer’ and ‘extra fine’ refer to default COMSOL mesh sizes & are applied to the bulk cavity material. A custom mesh is defined for the mirror (see eg Figure 42)

Boundary conditions (geometry)	Mirror mesh elements /1000	Bulk mesh	DOF /1000	Null cut [mm]	dl/l (1e-13) [/N]
¼ Symmetry Model	15	normal	73	<8.5	180
¼ Symmetry Model	39	normal	199	9.06	72
¼ Symmetry Model	83	normal	459	9.22	42
¼ Symmetry Model	143	normal	849	9.3	31
¼ Symmetry Model	20	fine	93	<8.5	182
¼ Symmetry Model	40	fine	204	9.08	66
¼ Symmetry Model	89	fine	457	9.26	37
¼ Symmetry Model	149	fine	875	9.31	29
¼ Symmetry Model	32	finer	147	8.54	174
¼ Symmetry Model	55	finer	271	9.09	65
¼ Symmetry Model	102	finer	538	9.25	38
¼ Symmetry Model	162	finer	932	9.31	29
¼ Symmetry Model	70	extra fine	297	8.64	148
¼ Symmetry Model	91	extra fine	412	9.11	63
¼ Symmetry Model	137	extra fine	681	9.26	37
¼ Symmetry Model	195	extra fine	1059	9.32	26

6.3 ANNULI GEOMETRY

Using the 5 cm cube model discussed above it is instructive to investigate the influence that the annuli (a component of the cube whose geometry can also be subject to change) has on the cut depth required for nulling the cube holding force. Here two parameters are considered: the annuli thickness, and internal radius. The standard annuli used in projects to date are 5 mm thick with 3.6 mm inner radii, and in this study both thickness and inner radii are varied individually to isolate the effect each has on the cut depth.

Figure 44 shows the fractional length change of the cavity measured for a sweep of annuli thicknesses using a ¼ symmetry FEA model. It is immediately clear that varying the annuli geometry provides an additional means of ‘tuning’ the geometry to produce an insensitivity to holding force, and indeed that multiple combinations of cube cut depth and annuli thickness could provide viable geometries. With the annuli inner radius held at 3.6 mm, the null point for this cavity (5 cm long, 1” mirrors, 7 mm cut depth) was found to exist with an annuli thickness of 3.98 mm.

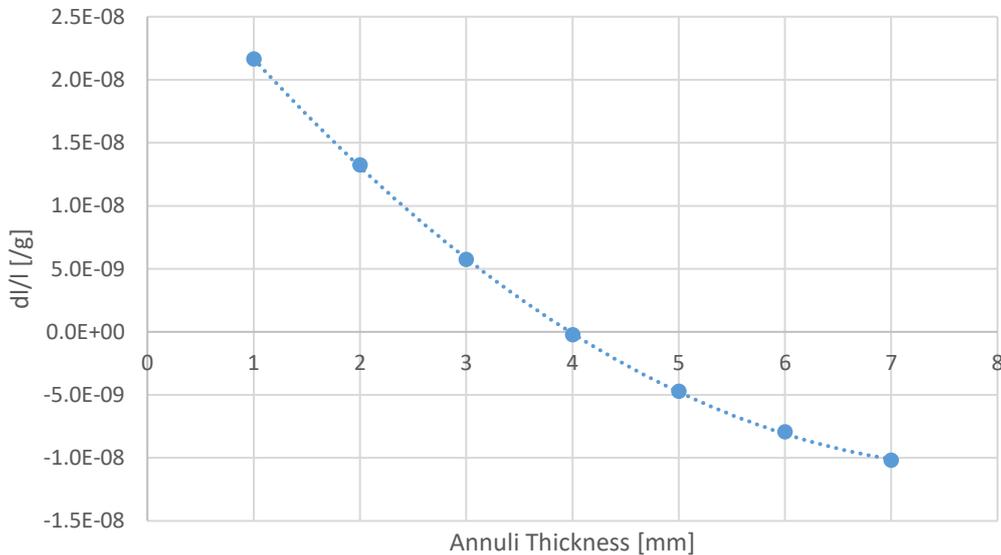


Figure 44 – Fractional length change vs. annuli thickness for a 5 cm cavity with 1” mirrors subjected to a 100 N compression force on the vertex

Figure 45 shows a similar plot to Figure 44, but here the annuli thicknesses were held at 5 mm, and the inner radius dimension was swept. By extrapolating a quadratic curve fitted to the measured data the null point was found to exist with an annuli inner radius of 5.58 mm.

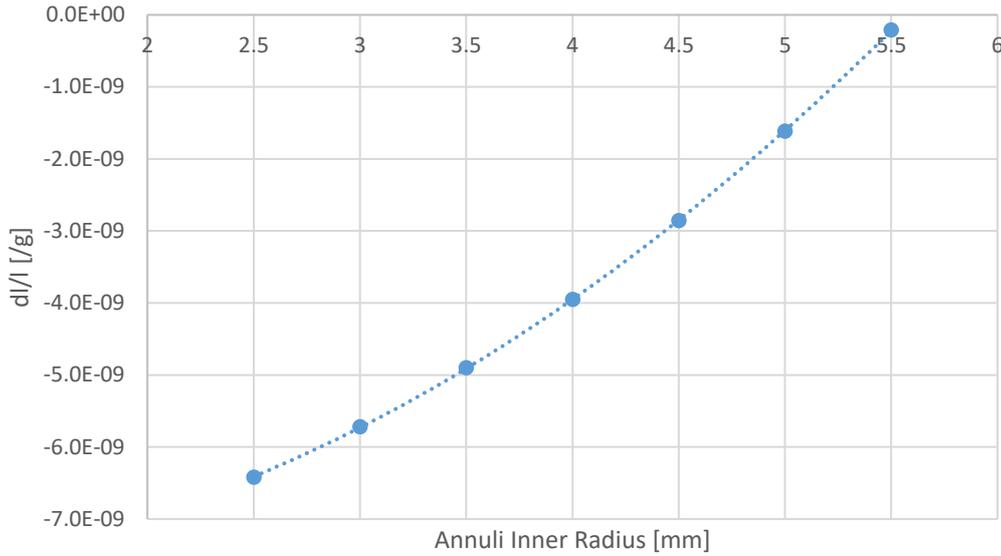


Figure 45 – Fractional length change vs. annuli inner radius for a 5 cm cavity with 1” mirrors subjected to a 100 N compression force on the vertex

6.4 3-CM CUBE (1” mirror with annuli)

Figure 46 shows the null point plot for a much smaller cube of 3 cm length assuming 1” mirrors with and without annuli of 5 mm thickness and varying inner annuli radii. It can be seen that varying the cut depth for this cavity, except in the case of an 8 mm inner radius, (within the limit available while ensuring the mirror is in full contact with the cube face) does not present a geometry with zero fractional length change. Reducing the size of the inner radii (*i.e.* increasing the amount of material the annuli possesses) shifts the fractional length change curve further away from the ‘no annuli’ case. In the case of the 8 mm inner radius annuli, a null cut depth of 4.75 mm exists.

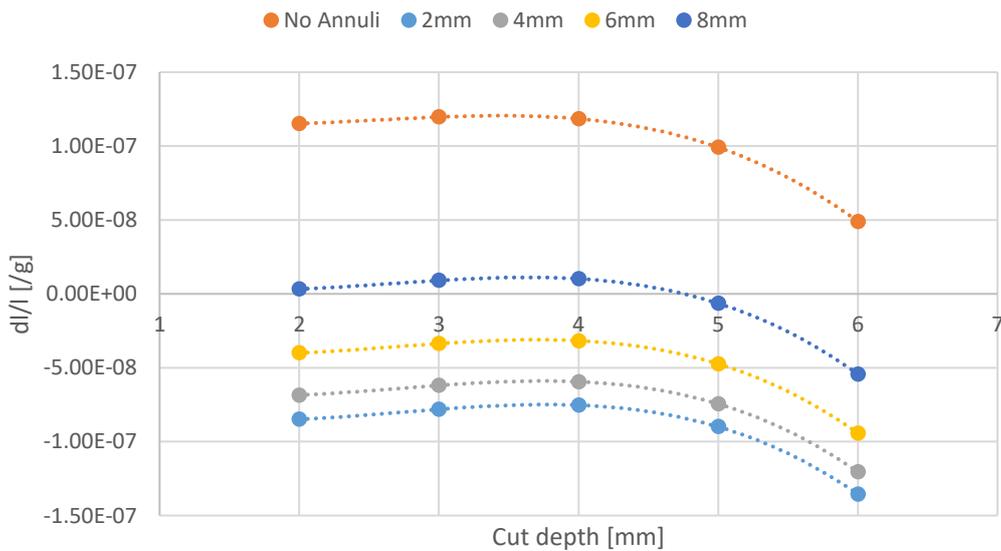


Figure 46 – Fractional length change vs. cut depth for a 3 cm cavity with 1” mirrors subjected to a 100 N compression force on the vertex. Annuli inner radius is varied while thickness is maintained at 5 mm.

Figure 47 shows a similar dataset, except now the annuli thickness has been reduced to 4 mm. It can be seen that the fractional length change curves have been shifted upwards towards the ‘no annuli’ case, suggesting that carefully balancing the annuli thickness and inner radius dimensions will allow a null cut depth to be found for this smaller 3 cm cube. In the case of the 8 mm inner radius annuli, a null cut depth now exists at 5.20 mm.

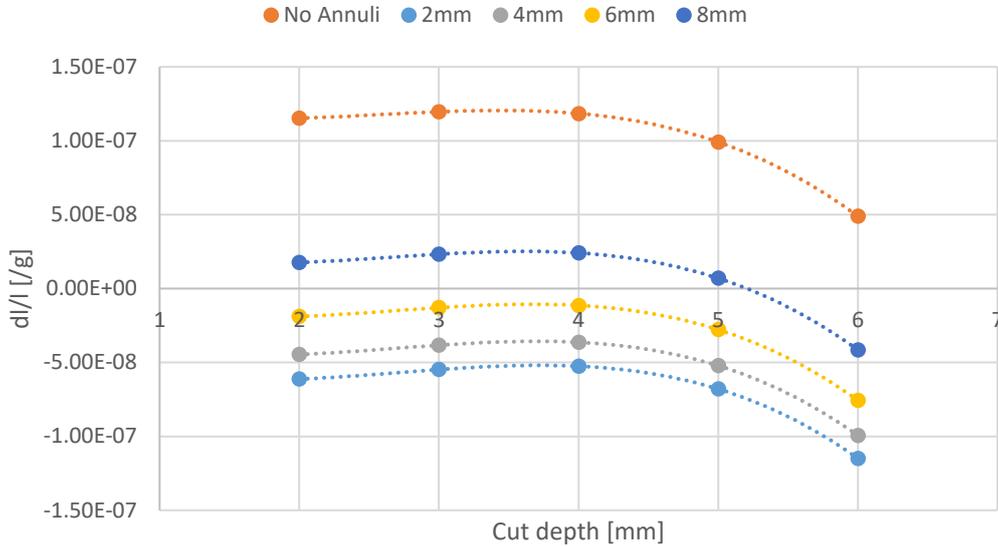


Figure 47 - Fractional length change vs. cut depth for a 3 cm cavity with 1" mirrors subjected to a 100 N compression force on the vertex. Annuli inner radius is varied while thickness is maintained at 4 mm.

6.5 4-CM CUBE (1" mirror with annuli)

Figure 48 and Figure 49 show null point plots for a small cube of 4 cm length and, while the magnitudes of the fractional length changes are reduced relative to the 3 cm cube, a similar pattern can be seen whereby tailoring the geometry of the annuli is required to enable a null point to exist.

Figure 48 shows the fractional length change vs. cut depth for cubes that have 4 mm thick annuli with varying inner radii dimensions, and Figure 49 shows the same plot but for 5 mm thick annuli. It can be seen that increasing the annuli inner radius or increasing the annuli thickness shifts the expected fractional length change down, and that for this cube with annuli of 4-5 mm thick, multiple null points can be found when the inner radii are between 4-6 mm. The null cut depths visible in Figure 48 are as follows: 9.73 mm when the annuli inner radius is 8 mm; 8.91 mm when the inner radius is 6 mm; and 6.73 mm or 7.43 mm when the inner radius is 4 mm.

NPLML – Commercial

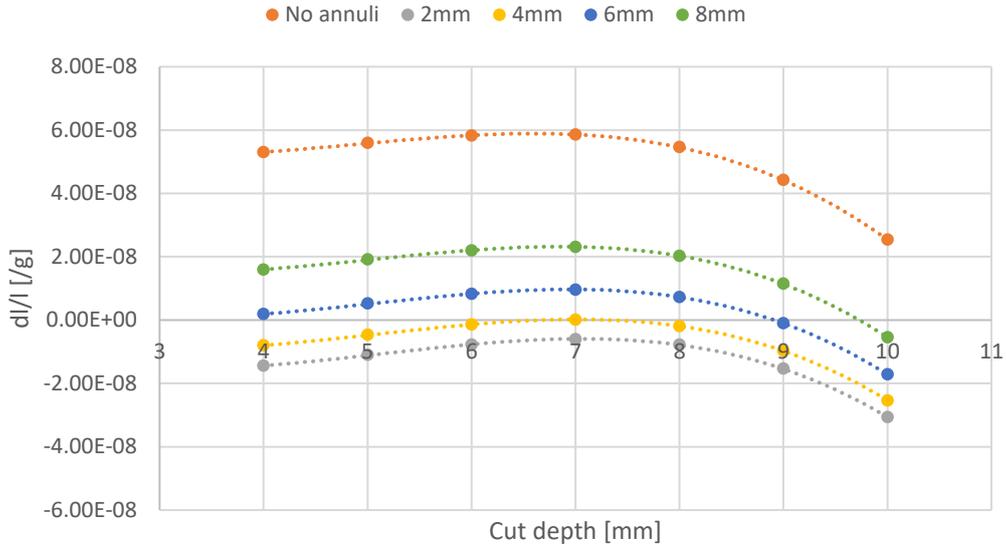


Figure 48 – Fractional length change vs. cut depth for a 4 cm cavity with 1" mirrors subjected to a 100 N compression force on the vertex. Annuli inner radius is varied while thickness is maintained at 4 mm.

Figure 49 shows the same plot as Figure 48 except that here the annuli thicknesses have been increased back to the original 5 mm. Here the plots of fractional length change have been shifted downwards with the increase in annuli thickness, and multiple null cut depths are clearly possible. The null cut depths visible in Figure 49 are as follows: 9.47 mm when the annuli inner radius is 8 mm; and 5.87 mm or 7.97 mm when the inner radius is 6 mm.

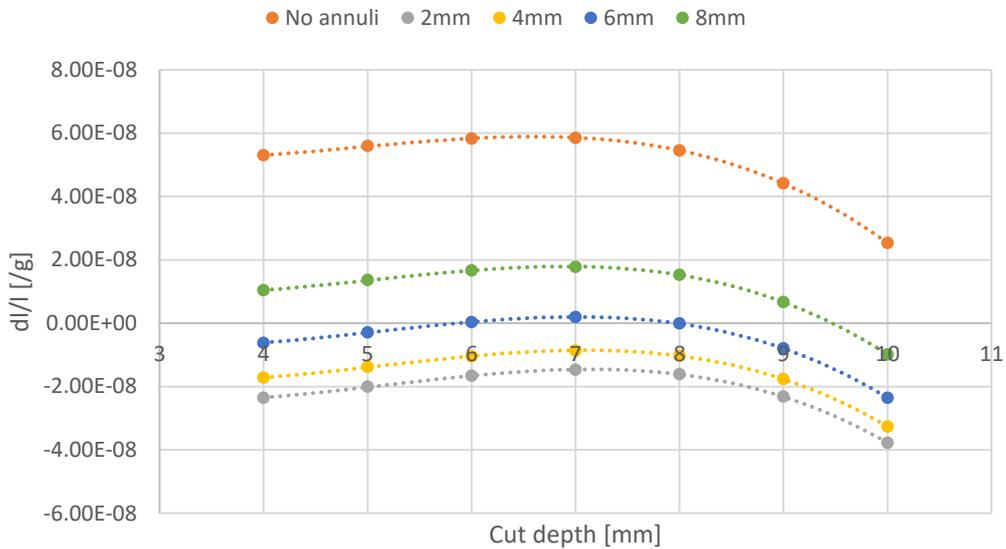


Figure 49 – Fractional length change vs. cut depth for a 4 cm cavity with 1" mirrors subjected to a 100 N compression force on the vertex. Annuli inner radius is varied with thickness maintained at 5 mm.

6.6 3-CM CUBE AND 4-CM CUBE (both with 1/2" mirrors)

Using 1" mirrors on smaller cubes limits the range of cut depths that can physically be achieved, and so investigating whether using smaller mirrors will also allow zero fractional length change cut depths to be found is valuable. Another common size for mirrors and annuli is 0.5" diameter and using these enables an additional millimetre or two of cut depth to be investigated. While the 1" diameter mirrors are 6.35 mm thick, these 0.5" diameter mirrors are 4 mm thick. Their smaller size will also reduce the resistance to deformation that the cube surfaces usually gain from having mirrors bonded to them, which may influence the optical axis fractional length change. It is also worth noting that although the mirror diameter is reduced by half, the approximate surface area of the ring that forms the optical contact is reduced to ~1/4 of that for the 1" mirrors.

Figure 50 shows the fractional length changes against cut depth for a 3 cm cube with 0.5" mirrors for various annuli inner radii sizes. For the 5 mm thick annuli investigated it can be seen that no suitable cut depth can be found with any inner radius dimension, with all measured fractional length changes being negative. Despite this, removing the annuli entirely shows positive fractional length changes, suggesting the annuli investigated are over-compensating.

Repeating this analysis with 4 mm thick annuli did not find a suitable cut depth either, suggesting that suitable annuli will be very thin. Future investigations into the use of very thin annuli may yield some suitable cut depths, but such annuli may not be feasible to manufacture. Alternatively, it may be possible to find a suitable cut depth using ULE mirror substrates (which may not need annuli) as opposed to fused silica mirror substrates, although the impact this would have on the cavity thermal noise floor would need to be assessed.

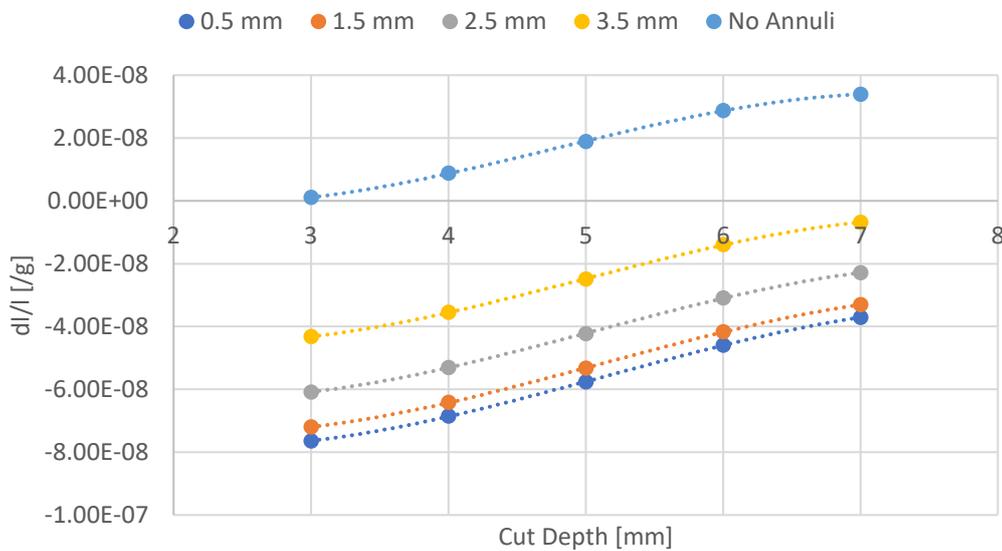


Figure 50 - Fractional length change vs. cut depth for a 3 cm cavity with 0.5" mirrors subjected to a 100 N compression force on the vertex. Annuli inner radius is varied while thickness is maintained at 5 mm.

Figure 51 shows a similar plot but for the larger 4 cm cube with 0.5" mirrors and annuli. Here it can be seen that null cut depths are possible for all sizes of annuli inner radii investigated. Assuming a cubic trendline the cut depths are: 4.74 mm with no annuli; 8.02 mm with 0.5 mm inner radii; 9.13 mm with 1.5 mm inner radii; 9.89 mm with 2.5 mm radii; and 10.23 mm with 3.5 mm inner radii.

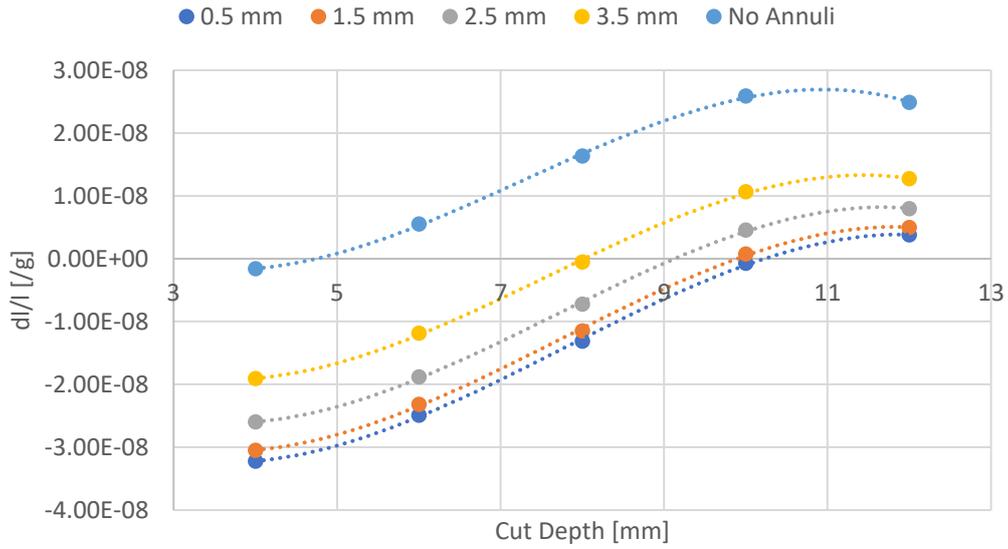


Figure 51 - Fractional length change vs. cut depth for a 4 cm cavity with 0.5" mirrors subjected to a 100 N compression force on the vertex. Annuli inner radius is varied while thickness is maintained at 5 mm.

The null cut depths visible in the plots within this section are summarised below in **Table 13**.

Table 13 - Summary of main null cut depths identified

Cube Length [mm]	Mirror and Annuli Outer Diameter [inch]	Annuli Thickness [mm]	Annuli Inner Radius [mm]	Null cut [mm]
3	1	5	8	4.75
3	1	4	8	5.20
4	1	4	4	6.73 or 7.54
4	1	4	6	8.91
4	1	4	8	9.73
4	1	5	6	5.87 or 7.97
4	1	5	8	9.47
4	0.5	N/A	N/A	4.74
4	0.5	5	0.5	8.02
4	0.5	5	1.5	9.13
4	0.5	5	2.5	9.89
4	0.5	5	3.5	10.23

6.7 EFFECTIVE CTE

The various geometric parameters that can be tuned in the design of a cubic cavity mean that the design space available to minimise fractional length changes is very large. While the studies described earlier by no means represent an exhaustive search of this design space, they give confidence that sufficient geometric tailoring routes exist across a range of parameter values that enable cut depths to be found for zero fractional length changes, even with smaller cube lengths. Using the annuli geometry and cut depths listed in **Table 13**, this section details the models used to assess the effective CTE that each one of these smaller cubes would exhibit.

As these ULE cavities have fused silica mirrors bonded to them (due to their lower thermal noise limits from such mirrors), an effective coefficient of thermal expansion must be evaluated for the optical axis as two dissimilar materials impact its length change. Length changes can be obtained from FEA models and the effective coefficient, α_{eff} , found from the expression of length change, dL , in terms of the original cube length, L , and a change in temperature, dT , as:

$$dL = L\alpha_{eff}dT \quad (11)$$

The CTEs for the ULE and fused silica are themselves functions of temperature and are approximated in the FEA models linear expressions. The instantaneous CTE for ULE

$$\alpha_{ULE_{inst}} = a_{ULE}(T_{cavity} - T_{ULE0}) \quad (12)$$

where the linear coefficient $a_{ULE} = 2.4 \times 10^{-9}K^{-2}$, T_{cavity} refers to the instantaneous temperature of the cavity, and T_{ULE0} refers to the temperature at which ULE is designed to have a zero CTE (this value can vary between 25-35 °C, and for the purposes of this modelling is assumed to be 30 °C). It is worth noting that the term a_{ULE} is positive when used in relation to length stability but would have a negative sign if it were related directly to frequency stability. The instantaneous CTE for fused silica

$$\alpha_{FS_{inst}} = 500 \times 10^{-9}K^{-1} + a_{FS}(T_{cavity} - T_0) \quad (13)$$

where the linear coefficient $a_{FS} = 2.2 \times 10^{-9}K^{-2}$ and $T_0 = 20$ °C. A quadratic polynomial is fitted to the resulting plots of fractional length change against temperature, which can be differentiated to give

$$\alpha_{eff} = \frac{dL}{LdT} \quad (14)$$

which allows the maximum thermal expansion and zero fractional length change CTEs to also be found.

6.7.1 3-CM CUBIC CAVITY

Table 14 shows the calculated temperatures where an effective zero CTE is found for the 3 cm cube, as well as the effective CTE for the cube length at 25 °C (a representative operating temperature). It can be seen that the effective zero CTE temperatures are very low – far below the anticipated operating temperature range for a space-based cubic cavity system. This suggests that for these smaller cubes the use of thicker annuli (or thinner mirrors if available) may yield more promising temperatures, but it was seen in Section 5 that finding suitable cut depths for this 3 cm cube was not straightforward, so time-consuming exhaustive searches of the geometric design space may be required. It is also worth noting that in a parallel project, assuming a higher T_{ULE0} increases the overall effective zero CTE temperature [37, 38]. In that case a zero CTE was reported 9.4 °C higher for a T_{ULE0} of 35 °C instead of 25 °C, but this is the upper limit of the range guaranteed by the ULE supplier.

Table 14 – Cube length effective thermal expansion coefficient for 3 cm ULE cavities with 1” fused silica mirrors and ULE annuli with 8 mm inner radii (ULE zero CTE assumed at 30 °C)

Cube Length [mm]	Annuli Thickness [mm]	Null cut depth [mm]	Effective Zero CTE Temperature [°C]	Effective CTE at 25 °C [K ⁻¹]
3	4	5.20	-13.6	-9.11E-8
3	5	4.75	-12.1	-8.76E-8

6.7.2 4-CM CUBIC CAVITY

Table 15 shows the CTE results for the 4 cm cube. As expected, it can be seen that thicker annuli and smaller radii favour the higher, more suitable zero CTE temperatures. For the smaller 0.5” mirrors, the zero CTE temperatures appear to be more favourable than those for the 1” mirrors, with some configurations offering effective zero CTEs temperatures >20 °C. This makes sense given that the 0.5” mirrors simply mean less fused silica (and hence less need for thermal compensation) in the system. The configurations with the best thermal performance require very small inner radii, so further study into these promising cases may be warranted to see if thicker annuli can enable more reasonable inner radii. For such a study it would be beneficial to better understand the limits on how thick the annuli can be made, and how small the inner radii can be cut.

Table 15 – Cube length effective thermal expansion coefficient for 4 cm ULE cavities with 1” and 0.5 ” fused silica mirrors and various ULE annuli (ULE zero CTE assumed at 30 °C)

Cube Length [cm]	Mirror and Annuli Outer Diameter [inch]	Annuli Thickness [mm]	Annuli Inner Radius [mm]	Null cut depth [mm]	Effective Zero CTE Temperature [°C]	Effective CTE at 25 °C [K ⁻¹]
4	1	4	4	6.73	10.1	-3.56E-8
4	1	4	6	8.91	3.4	-5.14E-8
4	1	4	8	9.73	-3.6	-6.78E-8
4	1	5	6	5.87	5.2	-4.70E-8
4	1	5	8	9.47	-2.5	-6.51E-8
4	0.5	5	0.5	8.02	22.8	-5.25e-9
4	0.5	5	1.5	9.13	21.1	-9.35e-9
4	0.5	5	2.5	9.89	17.8	-1.73e-8
4	0.5	5	3.5	10.23	13.3	-2.77e-8

6.8 SUMMARY OF SECTIONS 5 AND 6

Within sections 5 and 6 above, a series of Finite Element Analyses have been conducted to investigate the support configurations and geometric considerations of the cubic cavity for the CCU project, clarifying several key behaviours pertaining to the cube supports. Firstly, an important finding was that increasing the cube support force does not appear to negatively affect acceleration sensitivities regardless of the cut depth tolerance. Further studies indicated that cube positional errors may be the main cause of high inertial sensitivities seen in previous prototypes, and so more modelling based on planned measurements of cube alignments would be valuable, although the geometric design space has too many parameters to fully explore with FEA methods.

Comparisons between Hertzian theory, simplified 2D axisymmetric and full 3D sphere-cube contact models have also been presented, with full 3D contact models used to explore the forces required to prevent loss of contact under a worst-case quasi-static load. While this has highlighted a need to experimentally validate the expected friction coefficient between the chosen mounting spheres and a ULE cube, it is shown that contact forces in the order of several hundred Newtons are likely sufficient to prevent loss of contact or optical axis rotations of >5 arcminutes after a single 90g translational shock event. Furthermore, the magnitude of

these forces is unlikely to pose a cube fracture risk either. Low rotational eigenfrequencies were identified suggesting a potential vulnerability of the Hertzian contact support configuration to high torques, but such torques are not expected to be encountered in vibrational qualification tests of these cavity systems, and the translational modal masses at these frequencies appeared negligible.

An alternative mounting arrangement utilising ball-vee grooves in the cube was also investigated, showing that these grooves make the inertial sensitivities worse to an extent that they may not be capable of matching the performance of the original cubic cavity system [1]. As such, further investigations into using this mounting configuration are not recommended, although studies into alternatives (e.g. cone grooves on the truncated corners) may be worth pursuing. Any further analysis into mounting arrangements should wait until a clearer understanding on the cube positioning tolerances is acquired.

Trade-off studies have also been conducted to investigate the use of smaller cube and mirror geometries. It was found that the fractional length change could be tuned to zero by identifying a suitable combination of inner radius and thickness for the compensating annuli. For the geometries investigated, the use of both thicker annuli and / or smaller annuli inner radii correlated with smaller null cut depths and higher zero CTE temperatures. Suitable cut depths were found for cubes of 3 cm and 4 cm length, but the resulting zero CTE temperatures were often much lower than the anticipated operating temperature ranges of the cubic cavity systems.

For the 4 cm cube, suitable cut depths could be found when using 1" and 0.5" mirrors, although the resulting thermal properties were most promising with 0.5" mirrors, where zero CTE temperatures could be found at >20 °C. This configuration should be the focus of future studies in systems looking to use smaller cubes. It is worth noting though that reducing the cube size from 5 cm to 4 cm will mean the thermal noise is expected to increase by a factor of $\sim 5/4$. For the 3 cm cube suitable cut depths could be found when using 1" mirrors, although the resulting zero CTE temperatures were not suitable. Results suggested that suitable cut depths should exist when using 0.5" mirrors, although the compensating annuli required to make these work may be thinner, and the feasibility of using such annuli would need to be assessed.

It is worth noting that the findings of these trade-off studies are only valid for the geometries investigated. With cube length as well as the thicknesses and diameters (inner and outer) of the mirrors and annuli all available to be tailored, the design space for the cubic cavities is very large. A wider range of cut depths and annuli thicknesses may be worth investigating for the 3 cm cubes, as well as the effect of an increased thermal noise limit from the smaller cubes; however, the nature of this design space is that the required studies may need to be exhaustive and thus time-consuming. Alternatively, it may be possible to find suitable cut depths for a 3 cm cube with ULE mirror substrates, although this will require further modelling and an assessment of the impact on the cavity thermal noise limit. For now, the 4 cm cube (with 0.5" fused silica mirrors) offers suitable cut depths and thermal performance and is recommended as the first configuration to try in more compact future designs.

While this modelling activity has clarified understanding of the cavity mechanical support behaviours and presented pathways for design upgrades, several additional considerations surrounding the cube and support materials warrant further study, including:

- Experimentally investigating the Hertzian fracture risk. While ULE strength values from literature suggest fracture risk is low, we propose using softer indenter materials (PEEK) which may change the expected fracture stresses

- Experimentally investigating the Hertzian contact friction. While models were run for friction coefficient values between 0.1 and 0.3, there is no clear indication what value would be suitable for the ULE-PEEK contact.
- Assessing the impact of material creep. Given that the typical target lifetimes of these cavity systems are in the order of several years, it is worth investigating the expected impact of material creep – particularly as the support spheres are polymeric – on the contact forces and Hertzian stresses over the lifetime of the system.

7 CCU DUAL AXIS CCU OPTO-ELECTRONIC PERFORMANCE

The aim of the original clock control unit (CCU) project was to develop a dual axis cavity system for the frequency stabilisation of lasers required for a Sr lattice clock, except for the clock laser at 698 nm. The CCU is designed for operation at 461 nm (and the sub-harmonic at 922 nm), 813 nm, 689 nm, 679 nm and 707 nm. The CCU mirror coatings are also designed to provide specified finesse at wavelengths required for a strontium ion ($^{88}\text{Sr}^+$) clock. These include the sub-harmonic of the cooling transition at 843.3 nm (then doubled to 422 nm) and the $^{88}\text{Sr}^+$ clock transition at 674 nm. For a $^{88}\text{Sr}^+$ optical clock, CCU operation assumes that clear-out sources at 1033 nm & 1092 nm can be provided by LEDs [6]. Currently, the 698 nm Sr lattice clock laser is locked to a separate ULE high finesse cavity. In this way, one axis of the cube operates at 698 nm clock wavelength & is used for drift control of the other wavelengths [39].

As outlined earlier, the CCU comprises a DACC with optically contacted mirrors held in a tetrahedral mounting configuration in its aluminium frame. The precise positioning of the cavity within the centre of the frame is critical for achieving low acceleration sensitivity of the cavity. For optimum cavity performance, we use here an ultra-low expansion (ULE) spacer (Corning 7973) selected to have a zero linear coefficient of thermal expansion (CTE) at a convenient temperature control point. Typically, a zero CTE temperature is requested in the region between 25°C and 35°C. In the case where the optically contacted mirror substrates are also manufactured from ULE, the assembled cavity will have a similar thermal tuning characteristic to the ULE supplied. However, a lower thermal noise limit is possible using fused silica mirror substrates [8], although this significantly alters the thermal tuning characteristic of the cavity. The addition of optically contacted ULE annuli to the reverse side of the mirrors [8] partially offsets this effect. Our 5 cm DACC comprises 2 pairs of 25.4 mm diameter, 1 m radius of curvature mirrors, each with an additional ULE thermal-compensating annulus. The use of two pairs of 25.4 mm mirrors with compensating rings makes the cavity design less symmetric than [1] but numerical modelling confirms it is still possible to null the cube sensitivity to mounting forces through the truncation of the cube vertices by a specific length. In this arrangement, all the wavelengths required for a strontium lattice clock, except for the 698 nm (Sr) or 674 nm (Sr^+) clock laser wavelengths, are stabilised using one cavity bore. An orthogonal bore to this cavity axis is used to determine changes in the cube dimensions using the clock laser.

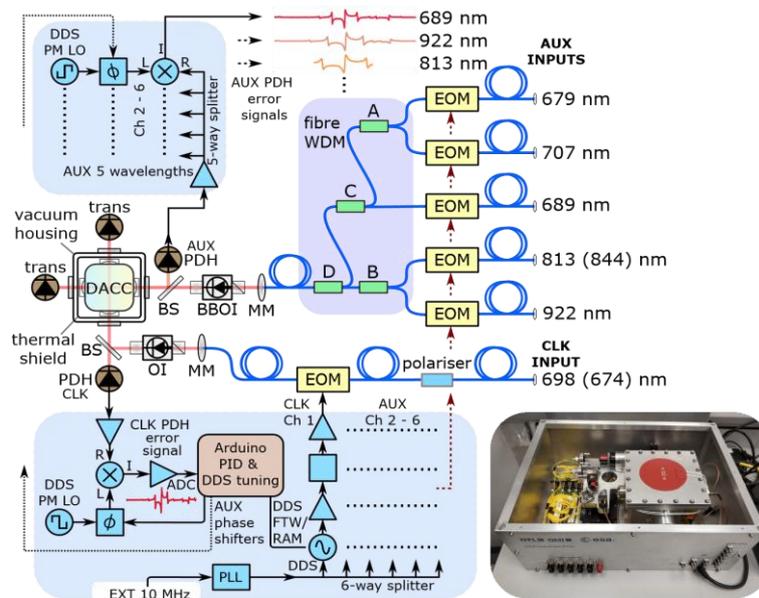


Figure 52: Schematic depiction of the CCU multi-wavelength stabilisation system exploiting the physical relation between orthogonal DACC axes to remove cavity drift. BS is a 50:50 beam-splitter and MM are fibre collimators that provide cavity mode matching. The complete optical assembly of the CCU, including the DACC, is contained within a 56 x 306 x 166 cm package, volume 23.2 l, total mass 15.9 kg. The electronics are housed in a separate 3U 19" unit.

The detailed optical schematic of the MW-CCU is shown in Figure 52. The different wavelengths required are multiplexed in via a fibre arrangement. The complete optical assembly of the CCU, including the DACC, is contained within a 56 x 306 x 166 cm package, volume 23.2 l, total mass 15.9 kg. This complete optics package is shown in Figure 5 the electronics are housed in a separate 3U 19" unit.

7.1 RECOMMENDATIONS FOR FUTURE DEVELOPMENT FROM THE FIRST PROJECT

This MW-CCU project follows from the first CCU project by addressing identified areas for CCU design improvement. Some of these areas of improvement have been covered in TN2, for example studies of different ways of mounting the cavity or the possibility of using different cube sizes. The question of designing a space qualifiable optics arm will be covered in a parallel ESA activity (a CCN to HSL2); improved digital electronics are being developed within TN4 of this study. Within this TN, we cover the remaining areas for improvement that include:

- It was recommended that the clock axis should have low-loss high finesse cavity mirrors (e.g. $F > 200,000$). These mirrors could either be dielectric mirrors designed for 698 nm or crystalline mirrors for 1397 nm operation. The advantage of operation at 1397 nm is that we would be able to take advantage of low-thermal noise crystalline mirror coatings [13] that demonstrate improved frequency stability over traditional dielectric mirrors [36].
- Further modelling of the cavity vibration insensitivity was recommended because a dual axis cavity with one axis used for clock laser frequency stabilisation needs to have optimum frequency insensitivity to acceleration. This is well understood in the case of the original single axis cavity but has not been so extensively modelled for the DACC. This modelling was discussed in TN2 of this project.
- Further experimental investigation of the ULE thermal expansion anisotropy was recommended, including a determination of the stability of the CTE expansion coefficients for the two orthogonal bores.

7.2 CCU CHARACTERISATION

This section summarises the results from the previous project on cavity finesse measurements, laser linewidth, frequency stability and characterisation of the cavity orthogonal drift (sections 5 to 7). In section 9.2 we describe the most extended recent cavity drift data, with more extended monitoring of the relative drift of the two orthogonal bores.

7.2.1 CAVITY FINESSE MEASUREMENTS

All four mirrors of the DACC have an identical high-reflectivity bespoke multi-layer dielectric coating (provided by FiveNine Optics, Boulder), with measured finesse above 10,000 at the critical 698 nm and 689 nm wavelengths, corresponding to a cavity fringe FWHM < 300 kHz. At 813 nm, where there is a more relaxed frequency stability requirement, a finesse (fringe FWHM) of 2,500 (1.2 MHz linewidth) was measured.

Testing was undertaken at four different wavelengths, 689 nm, 698 nm, 813 nm and 922 nm. All lasers except the clock laser are multiplexed in fibre to propagate along a single ‘auxiliary’ cavity axis with the clock laser fed to the orthogonal ‘clock’ axis (Figure 52). Both axes accommodate a free-space optical isolator (OI), which is broadband in the auxiliary axis (a broadband optical isolator, BBOI), to prevent parasitic etalons forming between the high-reflectivity cavity mirrors and upstream optical elements. Figure shows the transmission of the auxiliary axis of the cavity when these wavelengths are injected simultaneously into four different channels of the CCU. Each channel includes an EOM to generate sidebands that shift the laser frequency into resonance with the cavity. In a clock application, this sideband frequency can be used to tune the lasers into resonance with the appropriate clock transition frequency [7]. Here the EOM RF drive frequencies produced in the electronics module are scanned such that the sidebands for the different wavelengths are resonant with the TEM₀₀ mode at different times, leading to clearly resolved transmission peaks for the different lasers. The varying widths of the peaks is a measure of the different finesse at each wavelength.

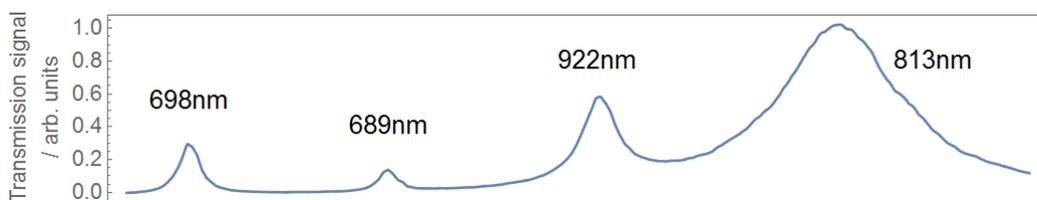


Figure 53. Transmission of the auxiliary axis of the cavity as the frequencies of four different injected lasers are scanned.

7.2.2 INITIAL CHARACTERISATION OF FREQUENCY STABILITY

Pound-Drever-Hall (PDH) error signals were measured by simultaneous scanning of lasers, and were observed to have no significant crosstalk (Figure 10 and repeated in Figure 54 below for ease of discussion). Losses through the multiplexed arm up to the cavity input are in the range of 18 to 27 dB, depending on wavelength, and 12 dB for the clock axis in the table below. Since the laser power requirements at the cavity are typically only a few tens of micro-watts, the losses are not problematic. The PDH discriminator determines the sensitivity of the error signal to cavity length fluctuations and is determined in part by the cavity finesse. Short term stability was assessed by a beat frequency comparison at selected wavelengths against independently stabilised lasers. We observed beat note linewidths of order a few kHz at

689 nm, which has the most stringent linewidth requirement, and a 1 s instability of ~ 200 Hz in the lock to a 698 nm sub-hertz clock laser.

Axis	Ch.	λ (nm)	IL (dB)	Finesse	FWHM (Hz)
Clock	1	698.4	12.3	11.3 k	266 k
Auxiliary	3	698.4	25.4	10.4 k	288 k
Auxiliary	3	689.5	27.1	8.8 k	342 k
Auxiliary	4	813.4	16.0	2.5 k	1200 k
Auxiliary	5	921.7	15.6	3.7 k	811 k

Table 16: CCU system properties. IL; insertion loss from input fibre to cavity input mirror, FWHM; full width half maximum of the Lorentzian cavity fringe-width

With modulation of the RF drive frequencies enabled, a series of independent PDH error signals are detected by a single photodiode in the auxiliary axis. These are separated during demodulation by use of discrete modulation frequencies, 8.900 MHz, 6.940 MHz, 16.625 MHz and 12.500 MHz for the locks at 689 nm, 698 nm, 813 nm and 922 nm respectively. These frequencies are chosen to minimise crosstalk while allowing sufficient clear bandwidth for the locking servo. Figure 4 (Figure 10 repeated here for convenience) shows the simultaneously recorded error signals when the EOM drive frequency of each channel is scanned across resonance. The frequency scans are set such that each laser passes through resonance at a similar time. With the laser powers balanced such that the error signals have a similar size, crosstalk between the different channels is inhibited. The secondary set of PDH features located about 32 MHz from the TEM₀₀ mode is due to another spatial mode of the cavity.

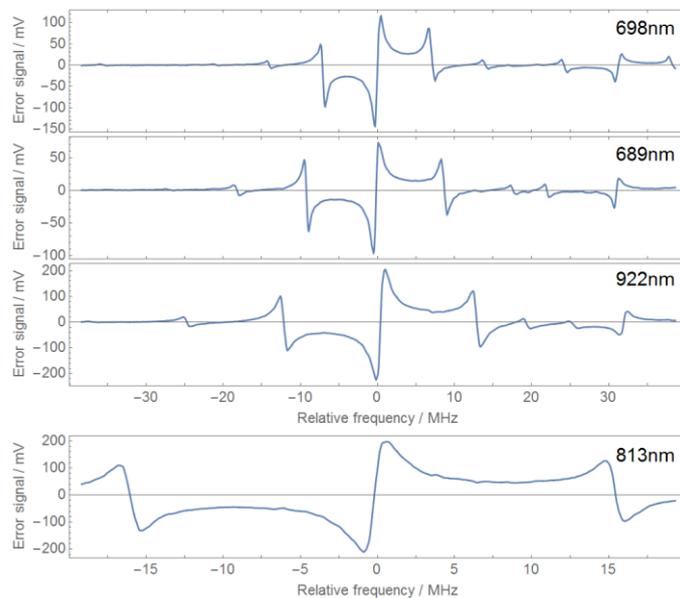


Figure 54. Simultaneously recorded PDH error signals from four different injected lasers.

Lasers are frequency referenced to the cavity fringes using the Pound-Drever-Hall (PDH) locking method in an electronic sideband modulation scheme [7, 40]. This provides a tuneable frequency offset f_0 between the laser carrier and cavity fringe by means of an optical sideband which is phase modulated to generate the PDH error signal. Analog Devices AD9910 direct digital synthesis (DDS) chips are employed to generate the phase modulated RF drive signals for each wavelength which are applied to the light through individual waveguide electro-optic

phase modulators (EOMs). An on-board state machine in the AD9910 allows for the rapid cycling of pre-programmed phase offset words to generate a phase modulated waveform, the period of which is made available externally and is subsequently processed through a programmable phase delay chip to provide the demodulation local oscillator (LO) for the PDH lock. The DDS chips are clocked by a 1 GHz phase-locked loop (PLL) with 10 MHz reference clock input. To bridge the 1.5 GHz half free-spectral-range (FSR) of the 5 cm DACC, each DDS is frequency doubled and amplified to provide f_0 up to 800 MHz. If a larger offset is required, up to the full half FSR, then the modulation depth is optimised for the second order sideband to which the cavity lock is made.

An Arduino microcontroller sets parameters of the DDS and phase shifting chips and is configured via a python script. The microcontroller also runs a low-bandwidth digital PDH lock which is implemented on the clock axis with feedback to f_0 via the DDS frequency tuning word (FTW). This provides a frequency record of the offset between the referenced (e.g. to an atomic clock transition) light which is coupled to the clock axis, and cavity fringe.

7.2.3 SUMMARY OF PREVIOUS LINEWIDTH AND STABILITY MEASUREMENTS

The key results from the design and pre-testing at NPL (prior to shipment to PTB for final evaluation) are:

- Measured cavity finesses (Table 16).
- Estimated design optical losses in the transmission to the cavity (Table 16) The temperature where the coefficient of thermal expansion is zero; 24°C and 26°C for the auxiliary and clock axes respectively.
- The thermal time constant (following a step change in temperature) was determined as ~8.5 hours.
- The electronics fitted within a 19" unit (Figure 8) and included the Arduino, clock circuitry, phase detectors and delay chips. The EOM drives were referenced to an external 10 MHz input.

Cavity drift ratios between the clock and orthogonal auxiliary axis were determined experimentally to allow drift compensation at all wavelengths that comfortably met requirements. **The results from PTB can be summarised as follows:**

- At 689 nm, the servo bandwidth was verified to be ~1.5 MHz from the position of a noise peak in the beats spectral output. The measured linewidth (determined via beats) was less than 3 kHz, which was limited by RF spectrum analyser bandwidth. The beats show a power spectral density of -62 dBc/Hz at 250 kHz offset. As this noise in the beat arises from both lasers, it is considered that the -65 dBc/Hz requirement is fulfilled.
- At 813 nm, the beats linewidth was 200 kHz and a noise level of -60 dBc/Hz was measured at an offset of 300 kHz.
- At 689 nm, a frequency drift of 0.5 kHz/hour was measured, within the 1 kHz requirement, using the measured relative drift between the clock and auxiliary axis. A drift rate suppression of a factor of 7.2 was demonstrated at 689 nm.
- At 813 nm, the drift over an hour was below 4 kHz, well within the 100 kHz specification.

Initial measurements of the isothermal drift of the two orthogonal bores were presented as part of the first CCU project. Here, we reproduce two plots for comparison with more recent data. Firstly, we studied the degree of correlation between the isothermal drift of each axis. Results of such a measurement are presented in Figure 55, recorded at a stable temperature of about 23°C. The chamber of the cavity was not actively temperature stabilised during the measurement, and residual temperature fluctuations are visible in the data. The inset shows

the result of plotting the frequency drifts of the two axes against each other, and a linear fit indicates a preliminary measurement of this gradient of 0.70. More recent and extended measurements are presented in the next section.

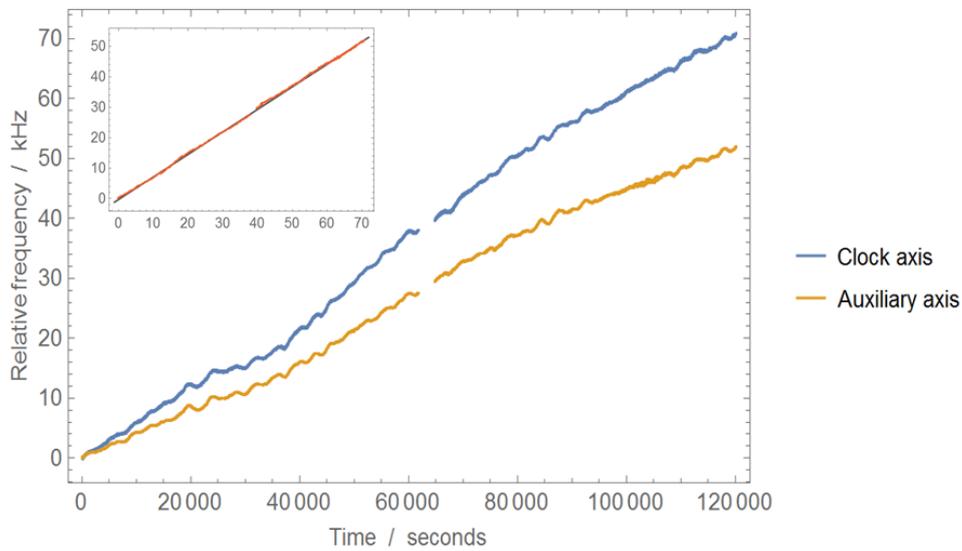


Figure 55. Measured isothermal frequency drift of the two cavity axes. Inset: plotted against each other, the drift of the two axes of the cavity are highly correlated. (Figure 12, repeated here for convenience)

Within the first CCU project, we also determined that the CTE crossover temperatures for both axes were different. These were approximately 24°C and 26°C for the auxiliary and clock axes respectively. This means that the CTE cannot be simultaneously minimised for both axes and complete compensation of thermally induced frequency fluctuations is not possible. The thermal time constant of the cavity within the vacuum housing was determined by applying a small frequency step of about 300 mK at about 23°C. The response of the mode frequencies of each cavity axis, shown in Figure 56, is well described by a simple exponential decay model and indicates a time constant for the system of about 8.5 hours.

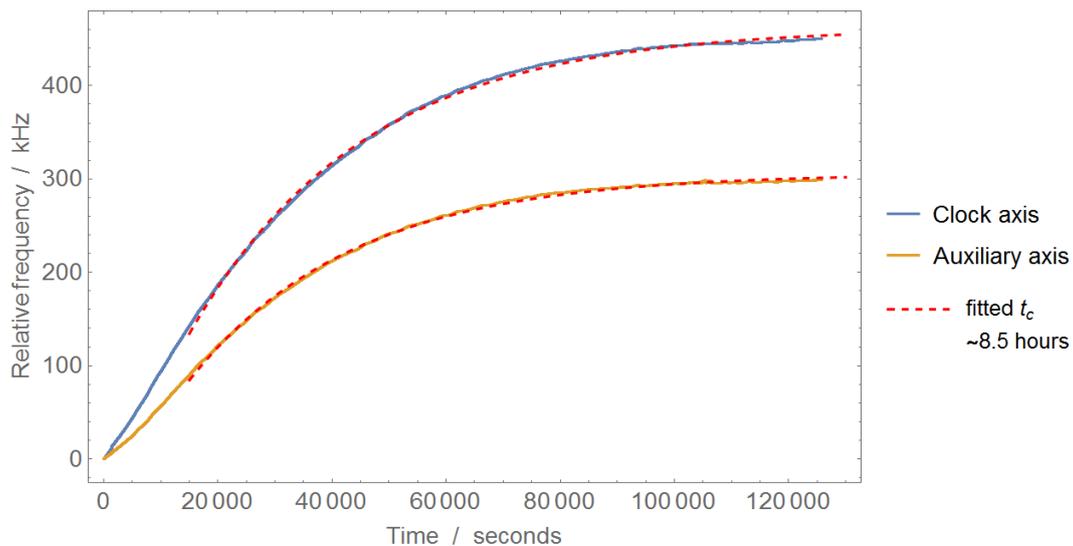


Figure 56. Measured frequency drift of the two cavity axes after a 300 mK temperature step. The fit indicates a cavity thermal response time of ~ 8.5 hours. (Figure 11, repeated here for convenience)

7.3 IMPROVED CHARACTERISATION OF THE CCU DRIFT RATES

The cubic cavity, mounted within its aluminium frame with single-layer heatshield, demonstrates a thermal time constant of 8.5 hours in response to a small temperature step. It is therefore not possible to control the cavity at a temperature where the thermal expansion is zero on both axes simultaneously.

In the analysis that follows, we envisage two use cases of the CCU:

1. The CCU could be used with a narrow linewidth laser that is not referenced to a transition in either Sr or Sr⁺. In this case, the differential drift of the two orthogonal arms is measured and used to predict the drift in the two orthogonal arms.
2. A clock laser is present and referenced to either Sr or Sr⁺, in which case the drift of one cavity bore can be directly measured and the bore drift used, together with a known characterised differential drift, to estimate the drift of the orthogonal bore.

Both scenarios assume that the differential drift is a constant of the ULE sample and remains constant over time. Data to date indicate that this is the case. However, as we will explain, the intended measurement scenario could affect the choice of ULE orientation – i.e. whether the two bores are both perpendicular to the growth axis or not.

In either measurement scenario, to compensate drift in the auxiliary axis we make use of the highly correlated creep of the orthogonal clock and auxiliary axes. The measured clock axis correction is independently scaled according to wavelength (as $r_l = \lambda_{clk}/\lambda_{aux}$) and fed-forward to f_0 of the corresponding auxiliary channel. An overall scaling factor, a , is included to account for the difference in rate of the correlated clock and auxiliary isothermal drift arising from anisotropy of the ULE glass spacer, such that the total scaling is $c = ar_l$. This scheme removes the need for an additional probe of the auxiliary axis for compensation of drift under certain conditions.

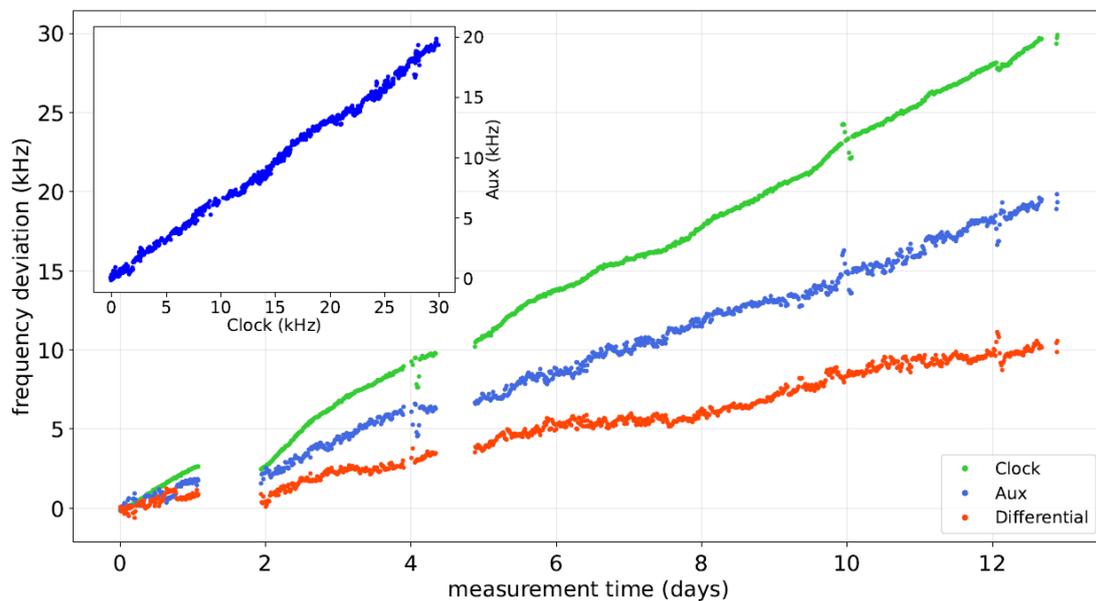


Figure 57: Measurement and compensation of the cavity drift by reference to a 698 nm ultra-stable laser over a two-week period. Linear drift rates of the clock and auxiliary axes were measured to be around 26 mHz/s and 17 mHz/s, respectively with correlation 0.643 (inset). Drift is compensated using both clock axis- and differential drift- feed-forward schemes.

To characterise and compensate drift of the DACC under conditions approaching isothermal creep, we stabilised a single ultra-stable 698 nm laser to both cavities and recorded the

required frequency corrections. Using the described Arduino-based digital lock to track the clock axis, and an additional synth-based analogue lock to the auxiliary axis, the corresponding frequency corrections, ν_{corr}^{clk} and ν_{corr}^{aux} respectively, were counted. The absolute frequencies of each DACC axis, ν_{abs}^{clk} and ν_{abs}^{aux} , were determined by the sums of the required frequency correction together with the absolute frequency of the 698 nm reference laser, ν_{ref}^{clk} , as measured by a hydrogen maser referenced optical frequency comb. The evolution of these frequencies was tracked over a two-week period, plotted in Figure 57, and the correlated relation extracted (inset). We find the scaling factor $a = 0.643$ (correlation coefficient $r = 0.996$), which can be exploited in a feed-forward scheme to compensate the auxiliary axis using a measure of the clock axis drift, shown in Figure 58 as ‘aux compensated’.

We believe that this asymmetry arises from anisotropy of the ULE glass, specifically that the properties along the material growth vector may differ to those of normal axes; this is currently the subject of further investigation. It may therefore be possible to choose a configuration of the DACC to provide a close to unity with cavities formed along normal axes of the glass spacer or maximise the difference by forming one cavity along the growth axis, as is the case here.

For the case where no clock laser is available, a value of a that is significantly different from unity is a potential advantage. In this case, this (assumed) fixed relation can be used as a reference in the removal of drift from both cavities without the need for an external optical reference. With a experimentally predetermined, by measuring only the differential drift of the orthogonal cavities in common view of a single laser, the absolute drift of each cavity can be extracted and removed by the following relations:

$$\nu_{diff} = \nu_{corr}^{clk} - \nu_{corr}^{aux} = \nu_{abs}^{clk} - \nu_{abs}^{aux}$$

$$\nu_{diff} = (1 - a) \nu_{abs}^{clk}$$

$$a = \nu_{abs}^{aux} / \nu_{abs}^{clk}$$

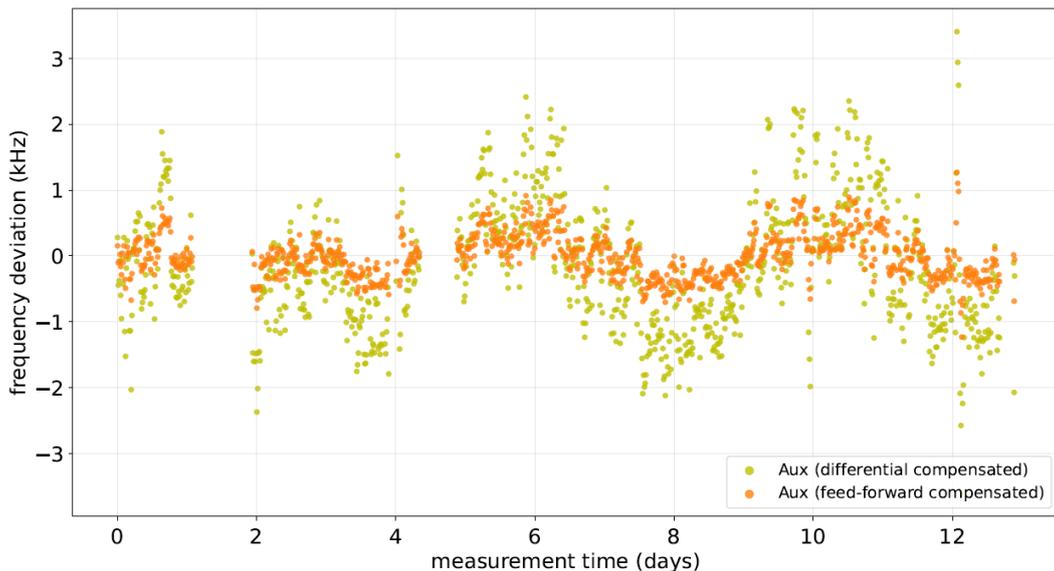


Figure 58: Residual frequency drift of laser locked to the auxiliary axis with drift compensation enabled.

The frequency deviation of both the clock and auxiliary axis is removed with an added instability of the measured differential drift inflated by the factors $1/(1 - a)$ and $a/(1 - a)$, respectively. This suggests that there could be an advantage in orienting the cavity bores to maximise the

difference in creep between both orthogonal axes. We applied this method to compensate the auxiliary axis drift, shown as Aux (differential compensated) in Figure 58.

The variation in linear thermal expansion of the two orthogonal bores with temperature could be used to an advantage in operating at a temperature that matches the relative thermal response of each axis to the ratio of isothermal creep between orthogonal cavities. This would mean that the correction would compensate both for isothermal drift and local thermal fluctuations. To understand how this could work, we write the frequency of a laser ν locked to a bore of the CCU as:

$$\nu = \nu_0 + Rt - \frac{1}{2}\alpha(T - T_0)^2$$

Here, R is the isothermal drift rate with time t and T_0 is the temperature at which the linear CTE is zero. The minus sign before α in the above equation arises because the cavity length is a minimum around $T = T_0$ and hence the frequency of a locked laser is at a maximum. If we consider the effect of thermal fluctuations $T_N(t)$ and operate at a temperature T_{set} ($T_N(t) \ll T_{set}$), then:

$$\nu = \nu_0 + Rt - \frac{1}{2}\alpha(T_{set} - T_0)^2 - \alpha(T_{set} - T_0)T_N(t)$$

As $(T_{set} - T_0)$ is a constant, we can redefine ν_0 to write:

$$\nu(t) = \nu_0 + Rt - \alpha(T_{set} - T_0)T_N(t)$$

It follows that there are two similar frequency correction terms ($\Delta\nu$) for the clock and auxiliary axis over the time interval Δt :

$$\begin{aligned} \Delta n^{clk} &= R^{clk} \Delta t - \alpha^{clk} (T_{set} - T_0^{clk}) T_N^{\dot{clk}}(t) \Delta t \\ \Delta n^{aux} &= R^{aux} \Delta t - \alpha^{aux} (T_{set} - T_0^{aux}) T_N^{\dot{aux}}(t) \Delta t \end{aligned}$$

Here, $T_N^{\dot{}}(t)$ is the time derivative of the thermal fluctuations. We have already measured experimentally the parameter $a \equiv R^{aux}/R^{clk}$. If the thermal fluctuations in the two bores are correlated, then we can write $T_N^{\dot{clk}}(t) \propto T_N^{\dot{aux}}(t)$. Expressed mathematically, we are suggesting operation at a set temperature T_{set} satisfying $a = b$, where b is defined as the ratio:

$$b = \frac{\alpha^{clk} (T_{set} - T_0^{clk}) T_N^{\dot{clk}}(t)}{\alpha^{aux} (T_{set} - T_0^{aux}) T_N^{\dot{aux}}(t)}$$

Since we have assumed that $T_N^{\dot{clk}}(t) \propto T_N^{\dot{aux}}(t)$, this ratio is independent of time t . If this condition is satisfied, then one coefficient ($\equiv a$) will allow correction of both the isothermal drift and random thermal fluctuations. We demonstrate in Figure 59 that a temperature should exist that satisfies the equation above. In the example below with T_0^{clk} and T_0^{aux} set to 24°C and 26°C, $a = b = 0.643$ is satisfied at 20.4°C. However, we note that this is not the temperature

at which the linear CTE is zero for either axis. Figure 59 assumes α^{clk} and α^{aux} are equal; the ratio $\alpha^{clk}/\alpha^{aux}$ determines this asymptotic value.

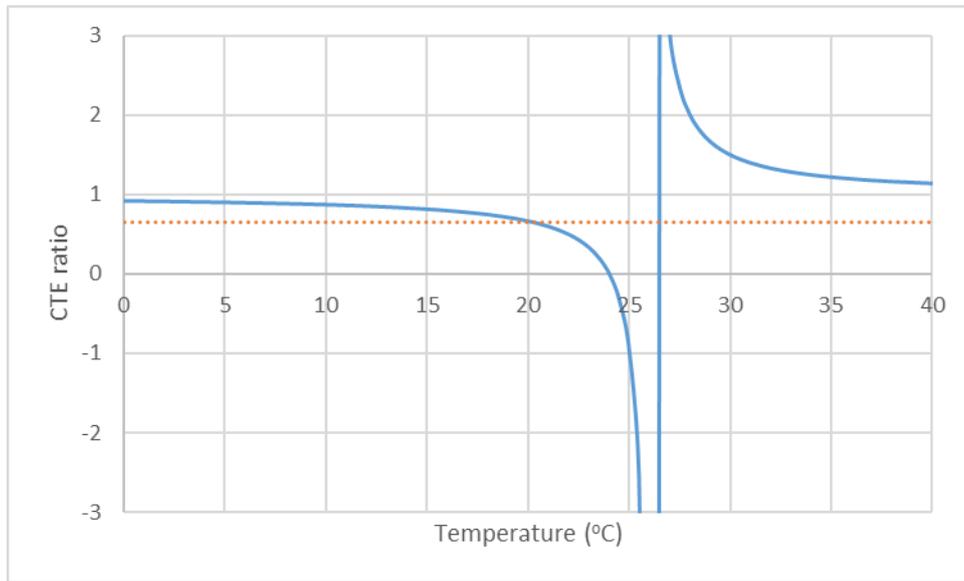


Figure 59: Ratio of the linear thermal expansion coefficients for the two orthogonal axes as a function of temperature. This plot can be used to find the temperature at which $b = a = 0.643$. The dotted orange line shows this value on the plot.

Finally, we reiterate that this scenario where $a = b \neq 1$ best suits a situation where the narrow linewidth laser is not referenced to an optical clock. If the CCU is primarily to be operated with either a Sr or Sr⁺ clock, then we suggest that a future cavity is oriented with both bores orthogonal to the growth axis. We then expect that the two cavity bores will have a similar thermal expansion coefficient and similar temperatures at which the linear CTE is zero.

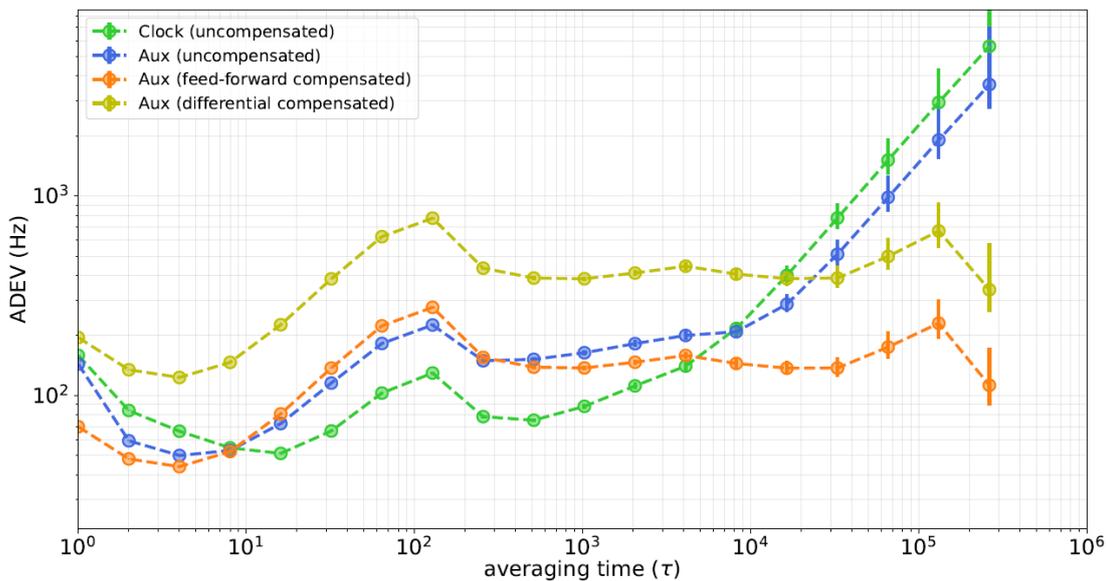


Figure 60: The cavity instability (overlapping ADEV) with & without drift compensation schemes applied.

The overlapping Allan deviation fractional frequency instabilities of the clock and auxiliary axes with and without drift compensation are calculated and plotted in Figure 60. A residual frequency deviation (instability) of < 1 kHz (< 200 Hz at 10⁵ s) is achieved for the compensated

auxiliary axis in the measured clock axis feed-forward scheme, which is inflated by a factor 2.8 (1.8) for the auxiliary (clock) axis where only the differential drift is used.

7.4 IMPROVED CHARACTERISATION OF THE CCU VIBRATION SENSITIVITY

The DC acceleration sensitivity of the DACC was measured simultaneously for orthogonal cavities through a series of inversions of the CCU module in x, y, and z directions, totalling 2g acceleration change per axis, while observing the required frequency correction to maintain lock to an ultra-stable reference laser at 698 nm. The results, shown in Table 17, are higher than [1] in some directions; we attribute this to slightly relaxed manufacturing dimensional tolerances, since low acceleration sensitivity was not specified for this project. Also, cavity mounting within the frame was via screws tightened to a much higher torque than previously to comply with high ESA shake and shock requirements. This might also have contributed to the increased acceleration sensitivity.

Cavity axis	Axial	Transverse (z)	Transverse (x,y)
Clock	$2.74 \times 10^{-11}/g$	$1.71 \times 10^{-11}/g$	$10.00 \times 10^{-11}/g$
Auxiliary	$0.24 \times 10^{-11}/g$	$3.00 \times 10^{-11}/g$	$9.65 \times 10^{-11}/g$

Table 17: DACC measured acceleration sensitivities in fractional frequency units per acceleration, g

Finally, we point out that if the CCU is used for frequency stabilisation of the clock laser, that optimised vibration insensitivity is more critical than for the originally planned configuration of the CCU where an independently reference clock laser was used. A combination of the further FEA detailed in TN2 and planned dimensional measurements of both an assembled cube and its position within a mounting frame are expected to steer future build tolerance specifications.

7.5 DESIGN DELTAS ON THE CCU CLOCK AXIS

In this section, we discuss the design requirements for a CCU that incorporates a high finesse cavity along one bore for a clock laser at 698 nm, 674 nm or sub-harmonics of these wavelengths. We also include discussion of options available for laser modulator, detector input beam mode matching at 1397 nm and frequency doubler to obtain 698 nm from 1397 nm.

7.5.1 CAVITY MIRROR OPTIONS

The choice available for mirror coatings at the clock wavelengths are:

- Dielectric coatings that would operate at both 698 nm (strontium lattice) and 674 nm (strontium ion)
- Crystalline mirror coatings are a lower thermal noise option if we operated at the respective sub-harmonic wavelengths of 1397 nm and 1348 nm.

For a 5-cm cubic cavity, a reduction of the thermal noise from $\sim 1 \times 10^{-15}$ to $\sim 5 \times 10^{-16}$ is expected in moving from dielectric to crystalline mirror coatings. Dielectric coatings are a well-established technology and suppliers have a proven track record of being able to produce low-loss optical coatings with design finesses of $\sim 200,000$ or more at both 674 nm and 698 nm. Crystalline mirror coatings are a more recent development although standard products are now available¹ at 1064 nm, 1397 nm, and 1550 nm. However, the coating bandwidth is much more limited, typically ± 20 nm. The relationship between cavity finesse (F) and mirror reflectivity (R) is:

¹ <https://www.thorlabs.com>

$$F = \frac{\pi}{1 - R}$$

At 1397 nm, the theoretical value of R is 99.9996%, giving $F = 785,000$. However, at 1348 nm (the 674 nm sub-harmonic), the calculated value of R is only 99.9877% yielding a finesse of only $\sim 26,000$. If operation at both 1397 nm and 1348 nm is critical, NPL could investigate with ThorLabs the feasibility of a bespoke coating with a ~ 50 nm bandwidth and centre wavelength of ~ 1372 nm.

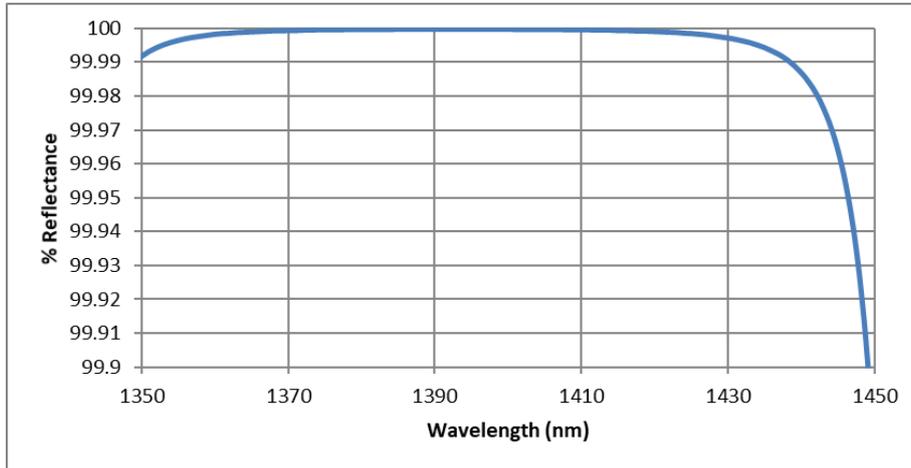


Figure 61: Theoretical reflectivity for crystalline mirror coatings (data downloaded from ThorLabs website 25th May 2021; ThorLabs part XM23P8)

7.5.2 LINEAR ABSORPTION LOSSES AT THE SUB-HARMONIC CLOCK WAVELENGTHS

A potential issue with operating in the near infrared is the effect of weak molecular linear absorption lines. These could either reduce the power at the cavity input or degrade the optical finesse. Degradation could arise either from weak absorption in residual gas in the vacuum chamber or water absorption in dielectric coatings. One supplier (AT Films) advertises a proprietary low water process to minimise absorption near 1390 nm. The linear absorption in a high finesse optical cavity will be enhanced by a factor $2F/\pi$ [41] over that observed single-pass over the same cavity length and so weak linear absorption lines could have a significant effect.

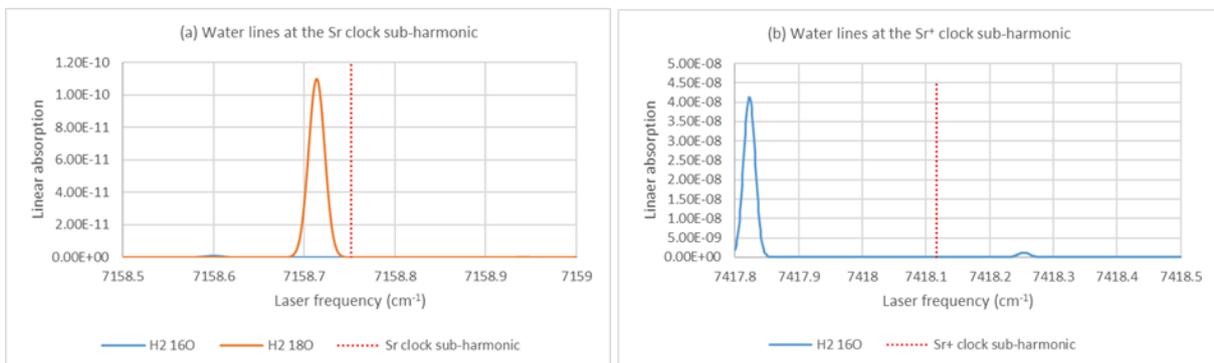


Figure 62: Linear water vapour absorption calculated for a 5-cm optical cavity of finesse 300,000 and evacuated to a pressure of 10^{-8} torr ($\sim \mu\text{Pa}$ level). A water vapour concentration of 1.15% is assumed;

at atmospheric pressure, this corresponds to a relative humidity of 50 % at 20 °C. The clock laser frequency sub-harmonics in cm^{-1} are shown as red dotted lines.

Water vapour is known to have strong lines in the ~ 1.4 mm region [42] and, in Figure we show the calculated absorption profiles for both the strontium lattice and strontium ion sub-harmonics. Calculations are for a typical cavity vacuum housing pressure of 10^{-8} torr (mPa level). The sub-harmonic frequencies were calculated from [43] for the strontium lattice clock and [4] for Sr^+ . We have also considered the effect of the linear absorption of methane and carbon dioxide. Of all these molecules, the strongest line is that of water (the principal isotopologue H_2^{16}O) a few GHz lower in frequency than the Sr^+ clock sub-harmonic (Figure 62). In free space at 50 % relative humidity and 20 °C, this line has a linear absorption of only $\sim 5.5 \times 10^{-6}/\text{cm}$. We therefore conclude that molecular absorption lines are not expected to cause a significant problem in this narrow spectral region.

7.5.3 FREQUENCY DOUBLING FROM 1397 NM TO 698 NM

Waveguided frequency doubling modules are commercially available from at least three suppliers worldwide at the present time. These are NTT², HC Photonics³ and Advr⁴; images of two of the three commercial products on offer are shown in Figure 63.

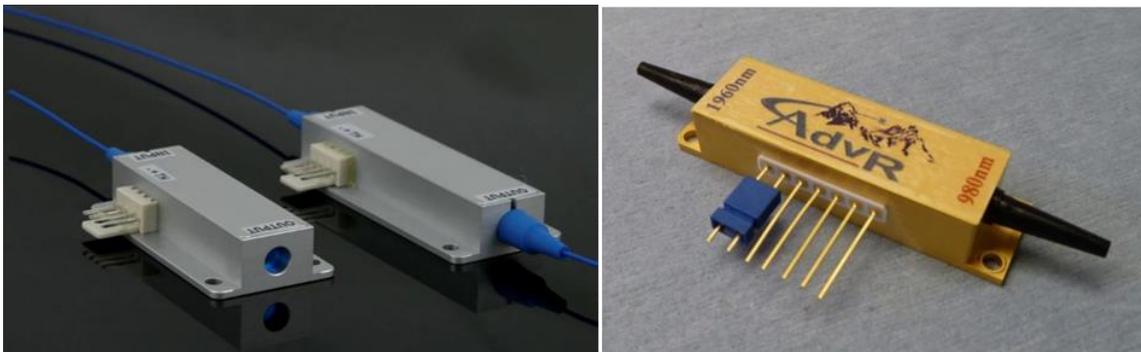


Figure 63: Commercially available frequency doubling modules available from HC Photonics (left) and Advr (right). The image of the HCP module was downloaded from <https://www.hcphotonics.com/ppln-mixers#3> and the Advr image from <https://advr-inc.com/waveguide-frequency-conversion/>

Confining the fundamental light within a waveguide creates a high beam intensity along the full crystal length; a combination of this high intensity and quasi phase matching at a specific temperature using periodic poling provides a means to generate the harmonic with high efficiency. The crystal doublers are provided with an integral Peltier and thermistor for temperature control at the appropriate phase matching temperature. The input is via a PM fibre and the output can be either into a PM fibre designed for the harmonic or into free space. Since the doubler is a waveguide, the output beam spatial mode quality into free space is close to TEM_{00} ; if the input fibre is slightly misaligned into the doubling crystal, output spatial TEM_{01} or higher order spatial modes can be observed at different temperatures.

As an example of the efficiency of these doublers, Figure shows the output power from an NTT frequency doubler as a function of the control thermistor resistance (\equiv temperature); in this case the doubler is designed to provide 461 nm from 922 nm. Only ~ 40 mW in free space is present just before the input fibre coupler and more than 0.8 mW is generated at the 461 nm harmonic output. An optical filter is used to block the fundamental at the output so that minimal sub-harmonic signal is observed at the output detector. With an estimated input fibre coupling

² https://www.ntt-electronics.com/en/products/gas_sensing/conversion_module.html

³ <https://www.hcphotonics.com/ppln-mixers#3>

⁴ <https://advr-inc.com/waveguide-frequency-conversion/>

efficiency estimated at 50%, the conversion efficiency is $\sim 0.8/20$, or $\sim 4\%$. As with all non-linear crystal conversion arrangements, the output power is proportional to the square of the fundamental power; the quoted efficiency, usually expressed as a $\%/W$, will increase linearly for low input powers at the fundamental.

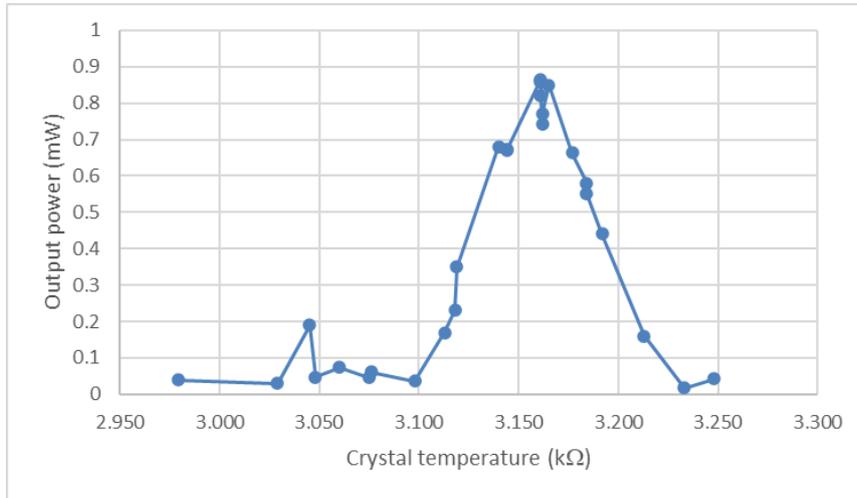


Figure 64: Output power in fibre from a frequency doubler converting 922 nm to 461 nm; the output power is shown as a function the control thermistor resistance.

7.6 DESIGN DELTAS FOR THE INPUT OPTICS

In this section, we review the input mode matching requirements at 1397 nm and discuss COTS component selection for the detectors, polarisation optics (for dielectric mirror coatings) and modulator.

7.6.1 Mode Matching into the cavity

The mode matching optics required are based on [44] which explains that mode matching from one waist size (w_1) to another (w_2) using a single lens of focal length f is only possible for a laser of wavelength l provided $f > f_0$ where the “characteristic length” f_0 is given by:

$$f_0 = \frac{\pi w_1 w_2}{l}$$

A straightforward conclusion from this equation is that, if we are to mode-match into the cavity over a short distance, both waist sizes need to be small. However, the waist size in the cavity (w_2) is determined by the cavity geometry. For a cavity of length $l = 5$ cm and mirrors that both have a radius of curvature R , the $1/e^2$ beam radius is given by [44]:

$$w_2^2 = \frac{l}{2\pi} \sqrt{l(2R - l)}$$

At $l = 1397$ nm and for $R = 1$ m this gives $w_2 = 263.5$ μm . Of the different collimators available commercially, part PAF2-2C from ThorLabs has a calculated waist of 380 μm positioned 73 mm from the lens. We note that these parameters are estimated by ThorLabs at a wavelength of 1550 nm and for SMF28 fibre. At shorter wavelengths, similar waist sizes are estimated by ThorLabs, but the focal distance is longer. To mode match to a waist size w_1 using lens of focal length f requires distances from the lens and focal positions of:

$$d_1 = f \pm \frac{w_1}{w_2} \sqrt{f^2 - f_0^2}$$

And

$$d_2 = f \pm \frac{w_2}{w_1} \sqrt{f^2 - f_0^2}$$

Either the positive or negative solutions to the above equations may be used. As an example, therefore, and using the positive solutions to the above equations, we have $d_1 = 221$ mm and $d_2 = 287$ mm if $f = 150$ mm. However, we note some corrections to this calculation that will be required:

- The design wavelength of these BK7 lenses is 587.6 nm (where the refractive index, $n = 1.5168$, whereas we are using them at 1397 nm ($n = 1.5025$). Using a thin lens approximation, the correction ($= (n_{588} - 1)/(n_{1397} - 1)$) will be ~ 1.0284 . Therefore, the lens focal length at 1397 nm will be 154 mm. This gives corrected distances of $d_1 = 230$ mm and $d_2 = 300$ mm.
- Although not considered here, a further small correction to the lens distance will arise because the cavity input mirror, which is plano-concave, will act as a weak lens.

For a future EQM optics arm, we note that lenses will require bespoke manufacture from radiation-hard glass. It should therefore be possible to request custom focal length lenses. We also note that an EQM system would require a space qualified collimator that is available from STI. This might not have the same optical properties as the ThorLabs device suggested here.

For $f = 125$ mm ($f = 129$ mm at 1397 nm), we could consider a comprise where the input beam size is calculated to be 0.22 mm (rather than the ideal of 0.26 mm). Then the distance from the collimator to the mode matching lens is 204 mm and the distance from the lens to the centre of the cavity 56 mm.

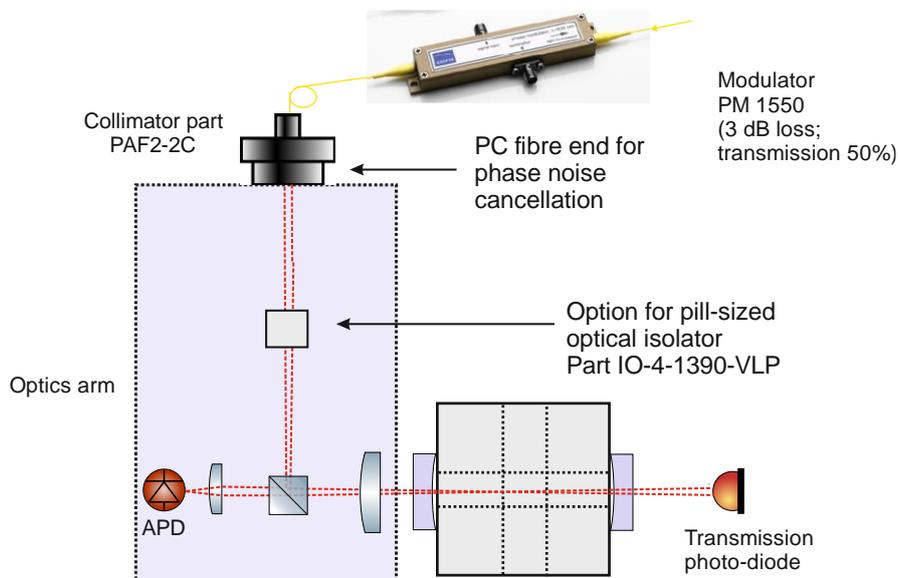


Figure 65: Input optics arm schematic for a cubic cavity clock axis. Example distances for different mode matching input lenses and collimators are given in the text.

We point out that, in Figure 65, the input optics lack the quarter wave plate in front of the cavity that might normally be expected for PDH locking. The quarter wave plate, together with a polarising beam-splitter cube makes a highly efficient optical means of recovering the reflected light from the cavity. However, this necessarily means that the cavity resonance is of circularly polarised light. Crystalline mirrors have a defined optic axis that is marked on the mirrors when supplied. These marks need to be aligned during cavity assembly and define the allowed directions for the linearly polarised input light. The cavity resonances are slightly different for the two different input polarisations allowed. Therefore, in Figure 65, the beam-splitter cube needs to be non-polarising and coated to allow (for example) 50% transmission and reflection.

7.6.2 COTS COMPONENTS

The following COTS components are available for use at 1397 nm:

- The standard ThorLabs quarter wave plate for 1310 nm has a retardation of 0.23 at 1397 nm (part WPQ05M-1310). A retardation closer to 0.25 could be obtained from an achromatic retardation plate (ThorLabs part AQWP10M-1600).
- For a polarising beamsplitter cube we could opt for ThorLabs part PBS104 that works from 1200 nm to 1600 nm (the part number is for a 10 mm cube).
- An InGaAs detector is required at this wavelength to monitor the reflected beam and a suitable COTS option is avalanche photodiode part APD310 (Menlo systems from ThorLabs) that works from 850 nm to 1650 nm.
- Either an InGaAs or germanium detector is also required to monitor the transmitted beam. Suitable detectors are ThorLabs part SM05PD6A (germanium) or SM05PD5A (InGaAs). Either detector could be operated photo-voltaically (i.e., with no bias) and should be sufficiently large that no focussing beam is required before the detector (Figure). These detectors also have an SM05 thread and so are direct replacements for the silicon equivalents (part SM05PD1A) that are used at 1064 nm.

The choice of electro-optic modulator will be steered by whether we opt for active RAM control [45]. If we do, then we would use a titanium in-diffused device that supports orthogonal linear polarisation and available from ixBlue and EOSpace. These modulators have a DC bias port that can be used to minimise RAM; temperature control of the modulator is also required. However, if RAM control is not required, then a proton exchange waveguide device would be better such as marketed by JenOptik. For these suppliers, the closest standard wavelength is at 1550 nm and performance at 1397 nm is not specified. JenOptik told us that the issue is that they do not have lasers at 1397 nm for testing, only at 1310 nm and 1550 nm. However, they do have waveguides available that, theoretically, should have a low insertion loss at 1397 nm (~3.5 dB) although they can only measure at 1310 nm, where the loss is ~4.5 dB. They have previously supplied modulators at this wavelength to customers including NPL but have no insertion loss data. We can presume that the situation is similar for ixBlue and EOSpace.

7.6.3 COMPACT OPTICAL ISOLATORS

We note that miniature optical isolators are available for this spectral region, for example part IO-4-1390-VLP from ThorLabs that has a maximum transmission and isolation specified for 1390 nm (Figure 66). At 1397 nm, we can expect a transmission of ~93% and isolation of ~34 dB. In addition to using this isolator in the optical arrangement for coupling light into the fibre, an isolator could also be required in the input optics arm.

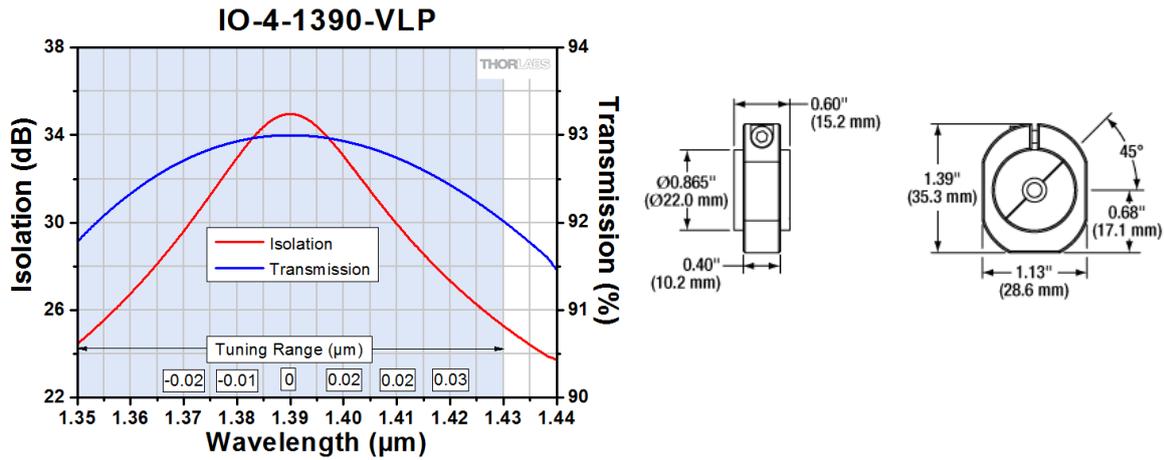


Figure 66: Transmission, optical isolation and dimensions of an optical isolator specified for 1390 nm (ThorLabs part IO-4-1390-VLP)

7.6.4 PHASE NOISE CANCELLATION

Depending on the length of fibre between the laser and the optical cavity, we note in this section that phase noise cancellation might be required [46]. A schematic of a suitable phase noise cancellation arrangement is shown in Figure 67. In this figure, light is reflected off the plane fibre end near the cavity (FC/PC connection) and double-passed through the acousto-optic modulator (AOM). The resulting beat at twice the AOM drive frequency, which includes twice the induced phase noise, is fed to a mixer and PID electronics to provide an error signal to the voltage-controlled oscillator (VCO). The drive to the AOM then includes a noise correction signal such that the frequency emitted from the fibre near the cavity has the same phase noise as near the laser head. This technique assumes that the noise induced by the fibre on the out and return paths are identical. Different reflectivities at the plane fibre end can be engineered by selecting different external semi-reflectors and aligning these back down the fibre. In this second arrangement, an FC/APC (angled) fibre end would be selected.

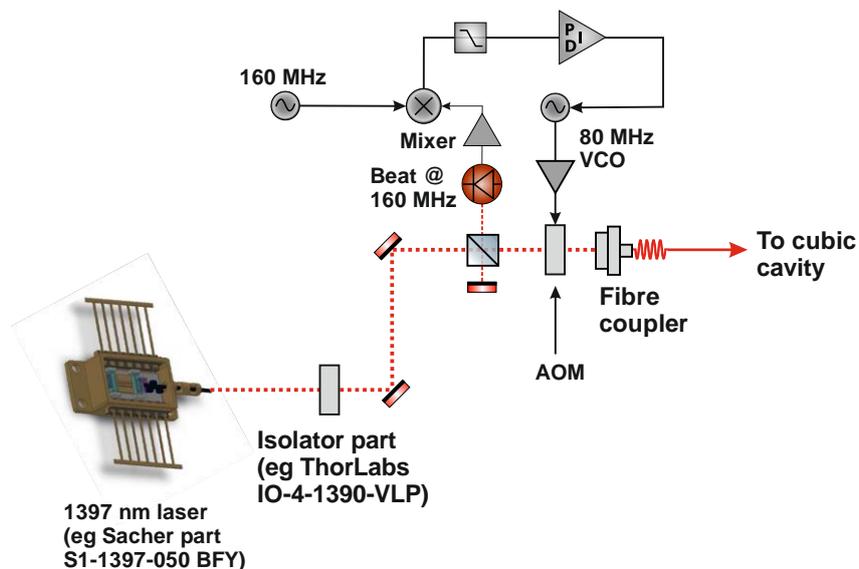


Figure 67: Schematic of laser launch and phase noise cancellation scheme. Light is reflected off the plane fibre end near the cavity (FC/PC connection) and double-passed back through the acousto-optic modulator (AOM). The resulting beat at twice the AOM drive frequency, which includes twice the induced phase noise, is fed to a mixer and PID electronics to provide an error signal to the voltage-controlled oscillator (VCO).

7.7 SUMMARY OF SECTION 7

Development of the DACC as a space compliant ultra-stable reference cavity is being pursued in parallel activities. A critical requirement is to increase the technology readiness level (TRL) of the cubic cavity and mounting arrangement. However further measurements and design changes will clearly also enhance the laboratory device. For example, by replacing the clock cavity axis mirrors with high-finesse alternatives, the same cavity could provide a reference for the ultra-stable clock laser with thermal noise limited instability at $\sim 1 \times 10^{-15}$ and below if crystalline mirror coatings are employed. The system would benefit from additional in-vacuum heat-shielding and improved knowledge of the CTE minimum temperatures for the two cavity arms to optimise the set temperature for the clock axis and better determine the operational temperature where the response is tuned to match the isothermal ratio. The long-term success of both cavity drift compensation schemes relies on the long-term stability of α which must be characterised over extended periods and with additional cavity systems. Measurements to date indicate its stability far surpass that of the linear drift rate and has potential to offer long-term Hz level absolute stability where sufficiently precise PDH locks are applied.

The results of the study so far highlight the successful development of a DACC-based compact clock control unit for a future space strontium lattice clock. The multi-band mirror coating demonstrated provides a suitable reference to satisfy all auxiliary laser stabilisation requirements in a single cavity axis. The device can easily be adapted to accommodate other systems based on eg. neutral Yb, Ca, Hg, Cd, Mg, or ions Yb⁺, Sr⁺, and Ca⁺. By determining the cavity drift using a 698 nm clock laser and also measuring the anisotropy of the drift we have shown that it is possible to provide a near drift-free output for all wavelengths, even without an ultra-stable optical reference, making the device a useful long-term stable frequency reference in its own right.

8 ELECTRONICS ENHANCEMENTS

This section details the design of electronics enhancements to the Multi-Wavelength Clock Control Unit (MW-CCU) laser frequency stabilisation and control. A digital fast Pound-Drever-Hall servo lock with MHz bandwidth has been designed and will be used for the stabilising all the cooling and auxiliary lasers in a strontium lattice clock.

The main design change to the electronics from the previous CCU implementation has been driven by the requirement to use the CCU also to stabilise the frequency of the clock laser, necessitating a higher bandwidth servo. A fully digital but lower bandwidth ~ 30 kHz (measured to be in closed-loop) PDH servo is being developed by NPL, required for other frequency stabilised laser applications. However, these electronics do not require frequency tuning to a strontium transition frequency which we implemented using the scheme in [7]. To demonstrate the feasibility of a fully digital arrangement, we will test a digital output PID stage with 1 MHz-plus servo bandwidth within this project.

Additional to the servo upgrade at 1397 nm, this redesign of the control electronics simplifies the electronics and makes larger frequency offsets more straightforward to generate. The previously used Analog Devices AD9910 DDS chips will be replaced by an RF DAC with integrated PLL. This allows the use of an offset sideband modulation scheme.

A description of the electronics system architecture design is presented first, followed by more detailed discussions of the component hardware and software design. Finally, the results of the verification of the designed digital PDH servo lock will be presented.

8.1 ELECTRONICS SYSTEM ARCHITECTURE

The MW-CCU control electronics have been designed as a modular system, comprising a ‘channel card’ for each laser frequency to be stabilised, and a backplane providing power supply rails and communications to the channel cards. A representation of the electronics system architecture is shown in Figure 68. Up to six channel cards can be supported in the system, and five of these will be allocated to the 461 nm, 679 nm, 707 nm, 813 nm and 689 nm lasers (R-01 in [20]).

Each channel card implements a configurable, fully digital PDH servo for locking the tuning actuator(s) of its associated laser, local DDS for the demodulation of the input PDH photodiode signal, and the generation of a tuneable offset sideband modulated signal for the EOM laser frequency tuning. Software configuration (pre-loaded or at run-time) allows the frequency, amplitude and phase of the DDS waveform used for demodulation and modulation to be set according to the frequency of the laser being stabilised. Similarly, the servo error setpoint, polarity, tune frequency and filter characteristics are also fully configurable (see Figure 75, which shows the functional layout of the channel card electronics). Two ADCs are provided for the PDH photodiode input and laser transmission strength inputs. The two 16-bit DACs will be used for the fast servo actuator output and optional direct modulation output. The 20-bit DACs are available to support a slow servo actuator output if required. The channel card electronics and software are described in more detail below.

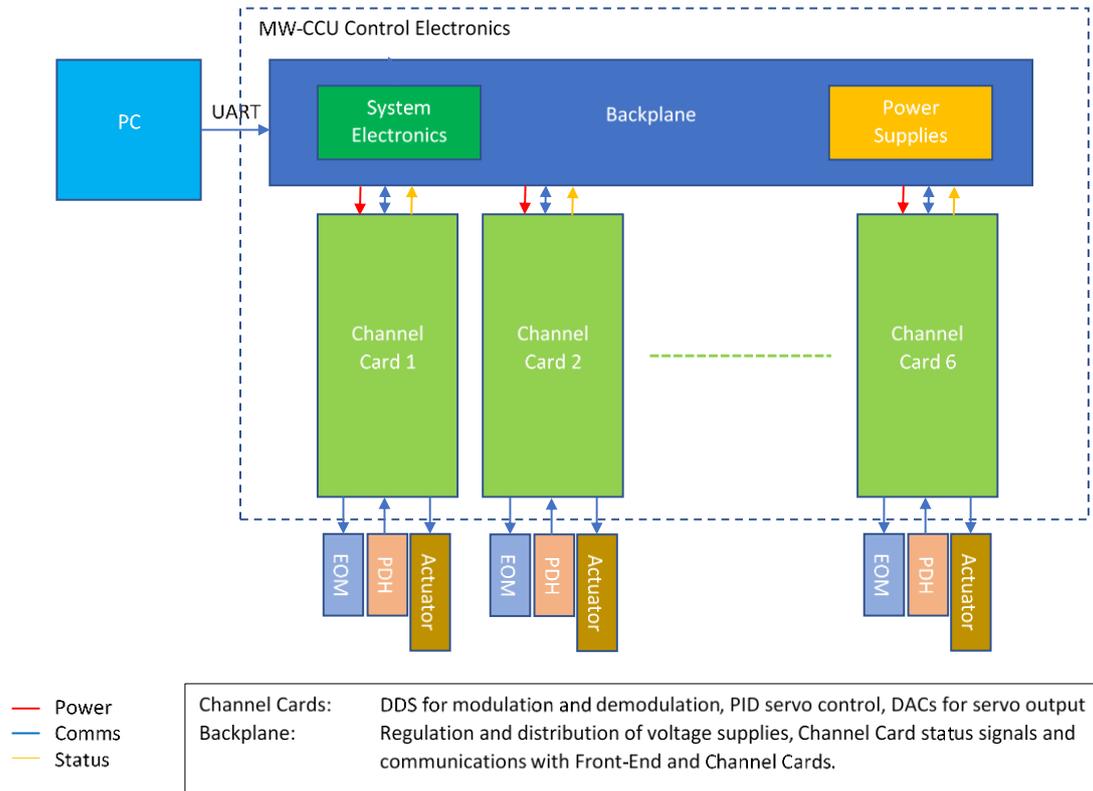


Figure 68 - MW-CCU Control Electronics Architecture

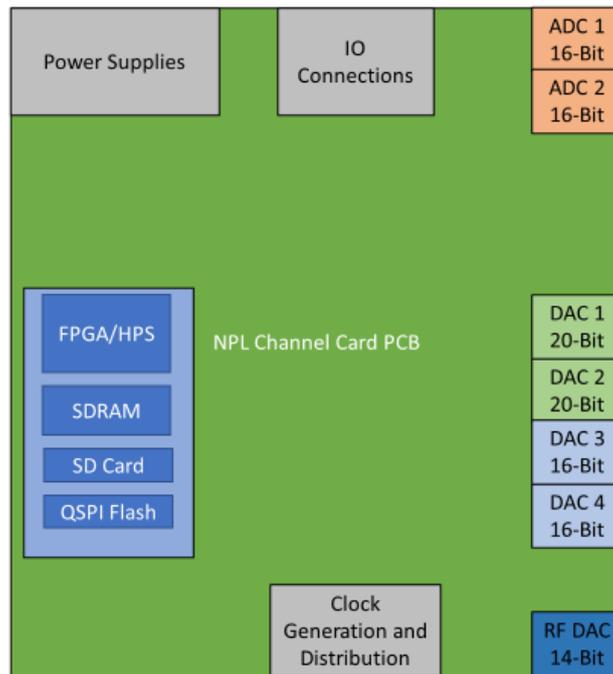


Figure 69 - Channel Card electronics layout

8.1.1 POWER SUPPLIES

The MW-CCU electronics are to be powered from a 12 V DC supply, which will be input to the backplane. From this supply, switched mode power supplies will be used to derive all voltage rails used by the electronics on the backplane and supplied to the channel cards.

A number of Analog Devices parts were identified as being potentially suitable for use in the MW-CCU power supply, namely the LTM4653, LTC3612, LTC3616 and the LTM4651. Of these the LTM4651 is an inverting regulator, which the other devices are non-inverting regulators. A brief summary and comparison of these components is shown below:

LTM4653 (Non-inverting ‘µModule’ – Low EMI)

- Around 80-95% efficient under ideal conditions
- No external inductor required
- Supports external clock sync if required (but not implemented in current design)

LTM4651 (Inverting ‘µModule’ – Low EMI)

- Around 80-95% efficient under ideal conditions
- No external inductor required
- Supports external clock sync if required (but not implemented in current design)

LTC3612 (Non-inverting DC/DC converter)

- Around 80-90% efficient under ideal conditions
- External inductor required
- Supports external clock sync if required (but not implemented in current design)

LTC3616 (Non-inverting DC/DC converter)

- Around 85-90% efficient under ideal conditions
- External inductor required
- Supports external clock sync if required (but not implemented in current design)

In addition, a RECOM RPM5.0-6.0 DC/DC converter has been selected to derive the 5 V supply rail from the 12 V input supply. This device has an efficiency of 90% and requires external inductors on the input and output for EMC filtering.

It should be noted that switching regulators have varying efficiencies, linked primarily to their internal design, as well as to their input voltage, output voltage, and applied load. At particularly high or low loads, the efficiency will drop off. Efficiency generally scales linearly with the difference in voltage across the device, with lower voltage drops offering higher efficiency.

The 3.3 V and 5 V supplies will be used by the microcontroller and USB to UART devices on the backplane respectively. All the derived supply rails, including the 3.3 V and 5 V rails will be provided to the channel card daughter boards via backplane connectors. Figure below shows the chosen power supply architecture to provide the required supply rails.

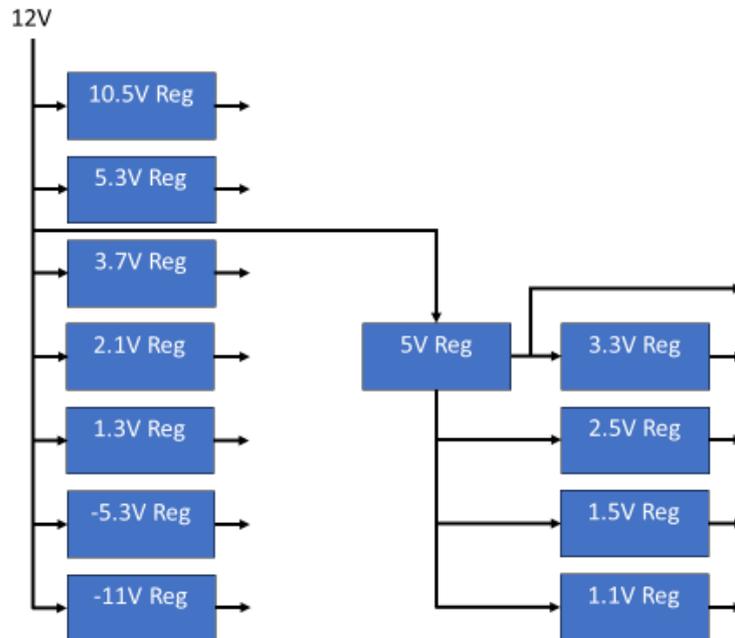


Figure 70 – Backplane Power Supply Architecture

The power consumption of the FPGA running on each of the channel cards will be a major contributor to the overall power consumption of the MW-CCU electronics. A full study of the system power consumption should be performed before PCB layout for manufacture is performed.

8.1.2 MICROCONTROLLER

The Atmel ATmega324 microcontroller is to be used on the backplane PCB, primarily for routing communication messages between the Front-End and individual channel cards. The ATmega324 is an 8-bit microcontroller with 32 KB of internal Flash memory, 1 KB of internal EEPROM storage and 2 KB of internal SRAM. The Flash memory has a minimum of 10,000 write/erase cycles, while the EEPROM has a minimum of 100,000 write/erase cycles. Data retention of both the EEPROM and Flash memory is at least 20 years.

The microcontroller will be clocked from an 11.0592 MHz crystal, a frequency that allows integer division for data communication down to 115200 bps and 9600 bps. Three serial USARTs are available of which one will be used for communication with the Front-End interface and a second will be used for communication between the backplane and individual channel cards. Passing data between two UARTs will add some latency to data transmission, however the messages being transmitted are for configuration and status reporting and as

such are not time-critical. Two LEDs are provided to indicate reception and transmission of data across the host interface from the Front-End.

The ATmega324 provides 39 General Purpose IO pins. Six of these have been assigned for the detection of channel cards being present in each of the available six channel card slots on the backplane. Each channel card provides 3.3V on one of its backplane connector pins. A 3.3V signal detected by the microcontroller for each slot will indicate a channel card is present. The Front-End interface will provide functionality to map each card slot to that slot's channel card function (that is, the laser frequency being stabilised by that channel).

A further six GPIO pins have been assigned by the same technique to allow each channel card to indicate its lock status to the backplane microcontroller, which can pass this information to the Front-End interface.

8.1.3 COMMUNICATION

The receive and transmit signals of the backplane micro-controller UART bus used for communicating with the channel cards, is shared to all channel card slots. Thus, any messages transmitted from the micro-controller will be received by the corresponding UARTs on every channel card. Messages from the Front-End interface via the backplane micro-controller will be addressed to specific channel cards, and to manage this a card position ID will need to be included in the message protocol and each channel card will also need to know its unique card position ID.

The user will know which channel card is providing frequency stabilisation for each laser frequency and configure the Front-End interface to map each channel card to a physical card slot. To support card slot identification each channel card connector on the backplane will set a unique 3-pin slot identifier. On each channel card these pins on its backplane connector will be wired to the card's microprocessor, allowing each channel card to read its card slot ID.

8.2 ELECTRONICS HARDWARE DESIGN – CHANNEL CARD

Much of the realisation of a digital frequency PDH servo lock sub-system for each of the laser wavelengths in the cubic cavity shares commonality with the digital PDH servo lock control electronics developed for related NPL projects on cavity stabilised lasers. For this reason, the MW-CCU "channel" control electronics design is based on these designs.

The key differences are:

- The MW-CCU modulation signal and servo bandwidth will be driven at higher frequencies to increase the servo bandwidth versus the requirements for optical clocks.
- The MW-CCU channel card electronics will also include an RF DDS and DAC device to perform offset sideband modulation.
- Designs for other NPL projects use a Intel Cyclone V FPGA and support devices on a Terasic DE10-Nano development board. The MW-CCU channel card electronics have been designed to incorporate the Cyclone V FPGA, SDRAM, non-volatile memory storage and support electronics on the channel card PCB directly.

Advantages of not using the DE10-Nano development board are that more FPGA GPIO pins are accessible for use, the power consumption of the channel card will be reduced as only required components will be fitted, and the overall size of the channel card PCB will be reduced.

8.2.1 FPGA AND MICROPROCESSOR

The Intel Cyclone V SE System on Chip (Soc) device, which contains both an FPGA and a hard processor system (HPS) within the same package will also be used on the CCU channel card. The HPS incorporates a dual core ARM Cortex-A9 micro-processor as well as 64kB of on-chip RAM, interface peripherals and a high speed HPS-FPGA interconnect. The HPS internal and peripheral clocks are derived from a 25 MHz oscillator input, while the FPGA will take a 50 MHz clock input.

The role of the micro-processor will be to manage communications from the Front-End user interface for parameter configuration, transmission of commands and lock status reporting. The micro-processor will also manage the programming of the FPGA, which must be performed on each power-up of the system. The HPS provides two UART interfaces and one of these will be used for communication with the Front-End, via the CCU backplane micro-processor. The other UART has been connected to a USB controller, to support direct communication to a channel card from the Front-End in the case where there may be a need to operate single channel card in a standalone mode (such as for test operations).

A 16 Kb EEPROM device that can be used to store any data that needs to be retained when the HPS or FPGA is powered off has been included in the design. Access to the EEPROM will be controlled by the HPS over an I2C bus. The HPS communicates with the FPGA via a shared bus interface, simplifying the transfer of data between the micro-processor and the FPGA.

All the DDS frequency generation and digital PDH locking control will be performed by the FPGA.

8.2.2 MEMORY DEVICES

The Cyclone V HPS has a 4 GB address space to support access to internal and peripheral registers plus program and data RAM. However, the HPS contains only 64 KB of on-chip RAM and 64KB of on-chip boot ROM. Non-volatile program code can be stored on a variety of external device types and must be transferred to RAM for execution by the micro-processor (although direct program code may be executed directly from a QSPI Flash device).

To support access to external SDRAM, the HPS provides an SDRAM controller along with the required address and data buses, control pins and clocks. A 512 MB DDR3L SDRAM chip (EtronTech EM6HE16EWAKG-10IH) has been selected to provide external RAM for the HPS and micro-processor. This exceeds the RAM requirements to support the microprocessor. The DDR3 SDRAM chip is powered from a 1.5V supply and will be clocked at 800 MHz. The chip's memory is internally organised into 8 banks of 32 Mb x 16-bit words.

The HPS supports the following external non-volatile memory types: SD card, Quad SPI Flash or NAND Flash through the provision of controllers and buses to interface with each of these memory types. Non-volatile memory is required to store the pre-loader and application software program code for the microprocessor and the FPGA firmware.

It was decided to provide two such memory devices in the design – an SD card slot and a QSPI Flash chip. QSPI Flash is the more secure of the two, as it cannot be physically removed from the PCB, unlike an SD card. However, a method to load the program code into the Flash memory device is needed before the software can be executed. For this reason, the SD card slot has been provided. Software can be loaded to an SD card to enable first execution of the pre-loader and application software. Once this software has been copied from the SD card into SDRAM is executing, the micro-processor can transfer the program code from the SD card to the QSPI Flash device. With the program code now present in the QSPI Flash device, at all subsequent PCB power on events the program code can be loaded into SDRAM from the QSPI Flash device.

A set of configuration pins on the Cyclone V control the non-volatile memory source to be used by the HPS. A switch has been provided in the design to change the configuration pin settings,

allowing selection of Flash memory source between the SD card or QSPI Flash device. The QSPI Flash chosen for the design is the 64 MB Cypress S25FL512SAG device, clocked at up to 108MHz with a minimum of 100,000 write-erase cycles and a minimum 20-year data retention.

8.2.3 CHANNEL CARDS POWER SUPPLIES

The backplane provides the source supplies for all power supply rails used on the channel cards. In all, there are eleven supply rails taken from the backplane, which are either regulated to derive lower voltage supply rails or used directly by devices in the channel cards. Figure 71 shows the backplane supply rails, the supply regulation employed on the channel card and the end use of each supply rail.

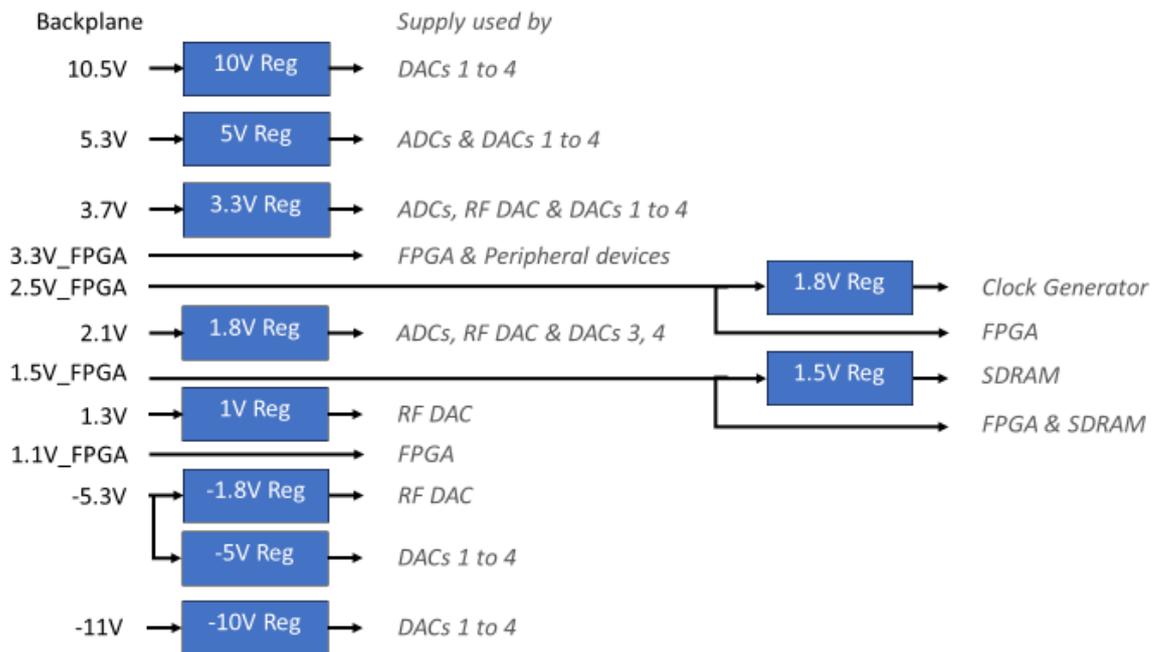


Figure 71 - Channel Card Power Supplies

The voltage regulators employed on the channel card are all linear regulators. The following device selections have been made:

The Analog Devices LT3045 ultra low noise linear regulator, with internal current limiting and thermal limiting protection will be used to derive the 10V, 5V, 3.3V, 1.8V and 1V supply rails from the 10.5V, 5.3V, 3.7V, 2.1V and 1.3V rails distributed from the backplane respectively. These supply rails are used to power the ADCs, DACs and RF DAC.

The Analog Devices LT3094 ultra low noise negative linear regulator, with internal current limiting and thermal limiting protection will be used to derive the -10V, -5V and -1.8V supply rails from the -11V and -5.3V rails distributed from the backplane. These supply rails are used to power the DACs and RF DAC.

For the ADC, DAC and RF DAC power regulation, optional precision voltage reference devices have been added to the hardware schematic should the high precision resistors used to set the regulation output voltage not be sufficiently accurate or stable.

The Analog Devices LT3080 low dropout regulator, with internal current limiting and thermal limiting protection will be used to derive the 1.8V supply rail required by the clock generator from the 2.5V supply rail distributed from the backplane.

Finally, the Diodes Incorporated AP2303 linear DDR bus termination regulator will be used to provide the 1.5V SDRAM termination voltage from the 1.5V supply rail distributed from the backplane.

Each of the linear regulators is enabled from a 'Power Good' signal, output by the switched mode power supply on the backplane that generated that regulator's input voltage.

8.2.4 CLOCK GENERATION AND DISTRIBUTION

This system is centred around a single input CDCE937 Low Jitter Clock Generator IC (hereafter referred to simply as 'the clock generator'). The input clock reference for the clock generator is an ECS 24 MHz, 10 ppm, SMD crystal. The clock generator has 7 clock outputs that can be individually programmed via an onboard EEPROM to required clock frequencies.

The clock generator will provide two main frequencies: 24 MHz will be provided to the Cyclone V HPS clock inputs, and 50 MHz will be supplied to the Cyclone V FPGA clock input. Some of the clock outputs of the CDCE937 clock generator are unused, and these will be turned off to reduce power consumption.

The default configuration of the clock generator will pass the 24 MHz input frequency through to all clock outputs. This will provide a valid clock frequency for the Cyclone V HPS, which will need to program the clock generator's internal EEPROM on first execution to provide 50MHz outputs for the FPGA. This will allow the clock generator to output the required clock frequencies on subsequent powering up of the PCB without further programming.

8.2.5 ADCS AND DACS

A key consideration when calculating the achievable lock bandwidth of the PDH servo lock system is the pipeline delay of each component. In an ADC/DAC context, the pipeline delay is the amount of time (typically in ns) taken for the ADC or DAC to sample or generate the appropriate voltage respectively. The pipeline delay is also a factor when considering the FPGA code, with extra processing and filtering taking more clock cycles to generate an output for the DAC, and therefore having a longer period of time.

An initial estimate for the maximum lock frequency is obtained by finding the reciprocal of the sum of all such delays in the system. As such, minimising each pipeline delay is important, particularly at higher lock bandwidths.

To limit the power consumption of the MW-CCU channel card, each of these devices and its supporting amplifiers can be completely physically isolated from the power rails when not in use. The specific designs will be explained in more detail in the following sections, but it should be noted that changes to the DACs in use should not be made while the channel card is in use.

8.2.5.1 ADCs

A study of available 16-bit ADCs was undertaken, with the shortlist shown in Table 17. Of these, the LTC2165 ADC (highlighted) was selected for NPL requirements, including the MW-CCU channel card. This device offers a 125 MSPS maximum sample rate with a 56 ns pipeline delay, at an average power of 194 mW per channel. This low pipeline delay provides the capacity for high bandwidth locking with the appropriate FPGA code.

IC	Resolution	Max Sample Rate (MSPS)	SNR (dB)	INL (LSB)	Latency (clk)	Pipeline Delay (ns)	Channels	Avg Pwr (mW)	Avg Pwr per Ch (mW)
LTC2165	16	125	76.8	2	7	56	1	194	194
LTC2195	16	125	76.8	2	7	56	2	431	215.5

Table 17 - ADC Selection Shortlist

8.2.5.2 DACs

Each MW-CCU channel card requires up to four DAC outputs, for the modulation output, fast servo output, optional slow servo output and an RF DAC for the offset sideband modulation.

As was the case for the ADCs, a study of available 16-bit and 20-bit DACs was undertaken, with the shortlist shown in Table . This study identified three potential options, of which two have been selected for the MW-CCU.

The AD9747 has been selected for output of the modulation signal and fast servo control. The AD9747 permits the use of one 16-bit parallel data bus to control two 16-bit DACs through use of an output selection pin. The power consumption per DAC channel is also fractionally lower than the LTC1668 at 155 mW, as are the INL and DNL figures. The maximum sample rate is 250 MSPS, and the estimated pipeline delay is 32 ns.

One pair of AD5791 20-bit DACs will be provided, offering high resolution outputs, at the expense of pipeline delay. This DAC will provide a locking bandwidth in the order required for a slow servo output but will not be sufficient to support the fast servo locking bandwidth requirements. This device was selected due to the proven heritage from the HSL2 project. The AD5791 has a maximum sample rate of 1 MSPS, operates at an average power of 126 mW, and has an estimated pipeline delay of 130 ns.

To allow flexibility in the choice of tuning range, a resistive divider has been provided at both the positive and negative reference inputs. By default, these are set to +5 V and 0 V, but can be adjusted to any voltage range between the +5 V and -5 V. This will allow the full DAC resolution to be applied to the appropriate tuning voltage range, up to a maximum of ± 5 V.

This reference voltage is then buffered using an ADA4004 operational amplifier and provided to the DAC. The DAC output is passed through an output amplifier which is configured to remove any offset from the input bias currents. Finally, the signal is passed through a DC-300 kHz low-pass filter to remove any high frequency imaging effects from the DAC being passed to the laser.

IC	Resolution	Max Sample Rate (MSPS)	INL (LSB)	DNL (LSB)	Latency (clk)	Pipeline Delay (ns)	Channels	Avg Pwr (mW)	Avg Pwr per Ch (mW)
AD5791	20	1	2	1		130	1	126	126
LTC1668	16	50	8	4	0	8	1	180	180
AD9747	16	250	0.1	0.05	8	32	2	310	155
AD9783	16	500	4	2	7	14	2	440	220
AD9726	16	400	1	0.5	4	10	1	400	400
AD768S	16	30	9	7	0	10	1	465	465

Table 18 - DAC Selection Shortlist (selected devices are highlighted)

A study of available RF DACs with integrated DDS was made and the shortlist shown in Table 19. The study selected the DAC38R84 device. It was decided not to use AD9910 from the original CCU project as the 1GSPS output rate is lower than the 1.5GSPS needed. It is noted that the apparent high-power consumption of the DAC38RF84 may be a concern.

IC	Resolution	Max Sample Rate (GSPS)	INL (LSB)	DNL (LSB)	Latency (clk)	Pipeline Delay (ns)	Channels	Avg Pwr (mW)	Avg Pwr per Ch (mW)
AD9164	16	12	2.7	1.7	1	26	1	2090	2090
AD9914	12	3.5	N/A	N/A	294	122	1	2237	2237
AD9910	14	1	1.5	0.8	91	101	1	715	715
DAC38RF84	14	9	4	3	916	611	1	3763	3763

Table 19 - RF DAC Selection Shortlist (selected device highlighted)

8.3 SOFTWARE AND FPGA FIRMWARE

8.3.1 FRONT-END INTERFACE

The Front-End interface has been developed to facilitate the changing of channel card configuration parameters from their default values set in the FPGA firmware of each card. The Front-End interface has been written in Python using Python’s built-in *tkinter* framework. This method allows rapid development of a GUI whose code can be run on many platforms without modification. The Interface application running on a PC communicates with the microcontroller on the backplane via a USB-Serial interface that appears as a COM port on the PC. The Python *PySerial* and *serial_asyncio* libraries are used to support asynchronous communication over this link.

On each power up of the CCU control electronics the FPGA on each channel card will operate from its set of configuration parameters contained within the FPGA firmware (and perhaps updated with parameters stored within the EEPROM device on each board). If no changes to these parameters are required there is no need to launch the Front-End user interface program.

The user interface program allows any configuration parameter on any channel card to be modified and this updated parameter will remain in use until the control electronics are powered off. If the updated parameter is to be used when the control electronics are next switched on, it will need to be set again using the Front-End interface program (or by sending the relevant command over the configuration parameter serial interface). An alternative would be to allow the EEPROM on each channel card to save the updated parameter if requested, so that the HPS can set this updated parameter in the FPGA on next power up. The Front-End user interface can save the configuration parameter value associated with any parameter ID, so that it doesn’t need to be re-entered manually each time the configuration parameter ID is selected in the program.

A simple and flexible approach has been decided for the interface design: to update a configuration parameter, the channel card number and the command number (or ‘identifier’) are selected from drop-down lists and the command’s default data is then displayed. This data can be modified, and the new value transmitted to the addressed channel card. The data must be entered as a 32-bit hexadecimal value. The available commands and their associated values are stored in JSON format within a configuration file specific to the MW-CCU project. A

'Parameter Management' screen is provided to add and/or modify and save commands and their values. Using this method commands can be added to the system without changing the Interface application code. The microcontroller code is described below together with the detail of the protocol used to communicate with the Front-End.

A facility is provided to write the current set of configuration parameters to an initialisation file specific to the selected channel card, which is used when programming the FPGA (so that the FPGA is preloaded with default values). This file is in the .mif (Memory Initialization File) format supported by the Cyclone V FPGA for initialising internal RAM.

A 'Help' screen provides the currently available commands that can be used.

Screenshots of the main user interface, parameter management and help pages are shown in Figure 72, Figure 73 and Figure 74 respectively. The existing interface application has not yet been updated to provide the channel card selection drop-down lists that will be used to allow parameter selection and modification specific to the selected channel card.

The 'Initiate Frequency Sweep' button shown on the main interface screen is provided to instruct an FPGA to run an automated frequency sweep. This will result in locking once the scan is complete and a peak in the transmission photodiode signal above the frequency scanning threshold has been detected and applied.

When the program is started the user first opens the COM port specifying the microcontroller interface. The user can then send commands in the main interface page while referencing the provided command list.

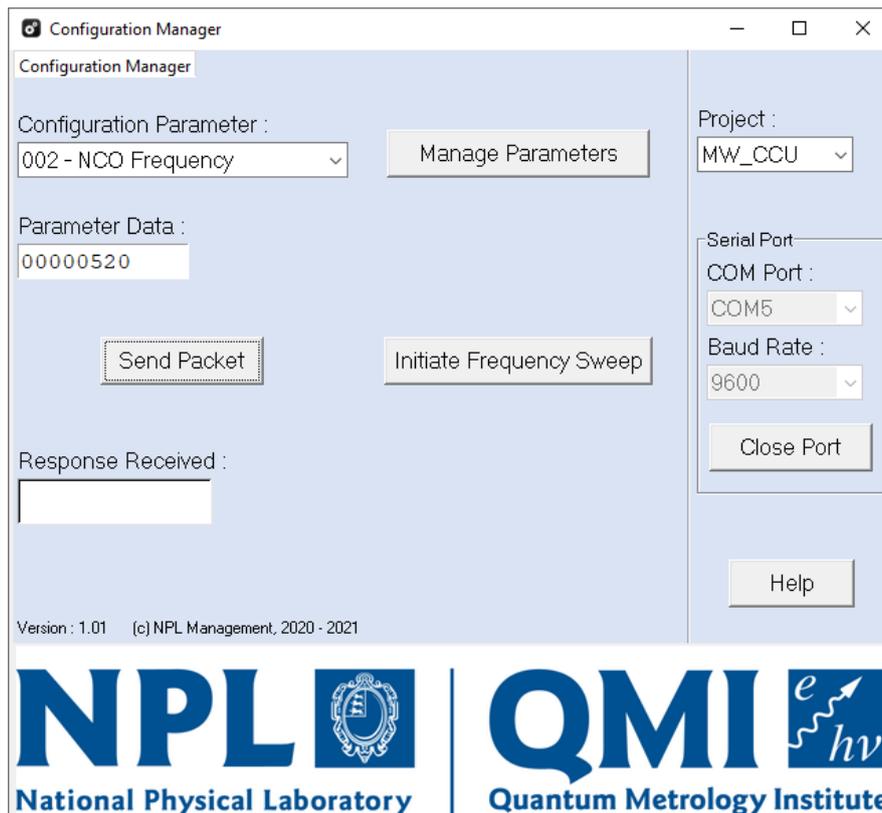


Figure 72 - Front-screen of the GUI interface. The user can send data to the device by selecting the identifier, modifying the data and clicking the 'Send Packet' button. Data received in the message response will then be displayed.

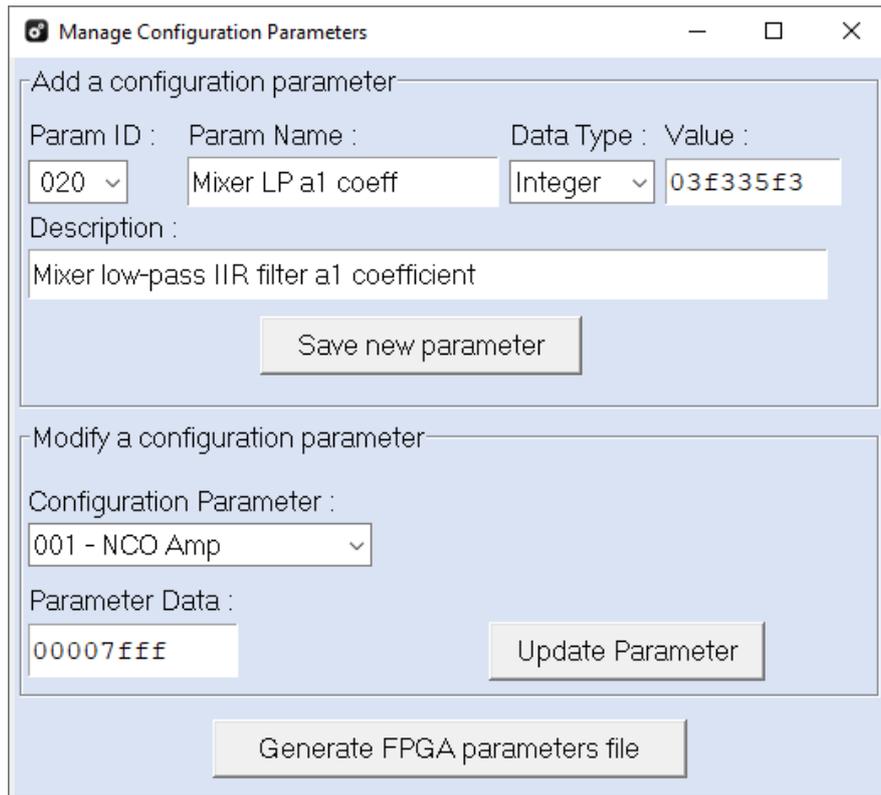


Figure 73 - Manage Configuration Parameters' screen of the GUI. New configuration parameters can be added and existing parameters modified from this screen. A button is provided to generate a file containing parameters to pre-load into the FPGA.

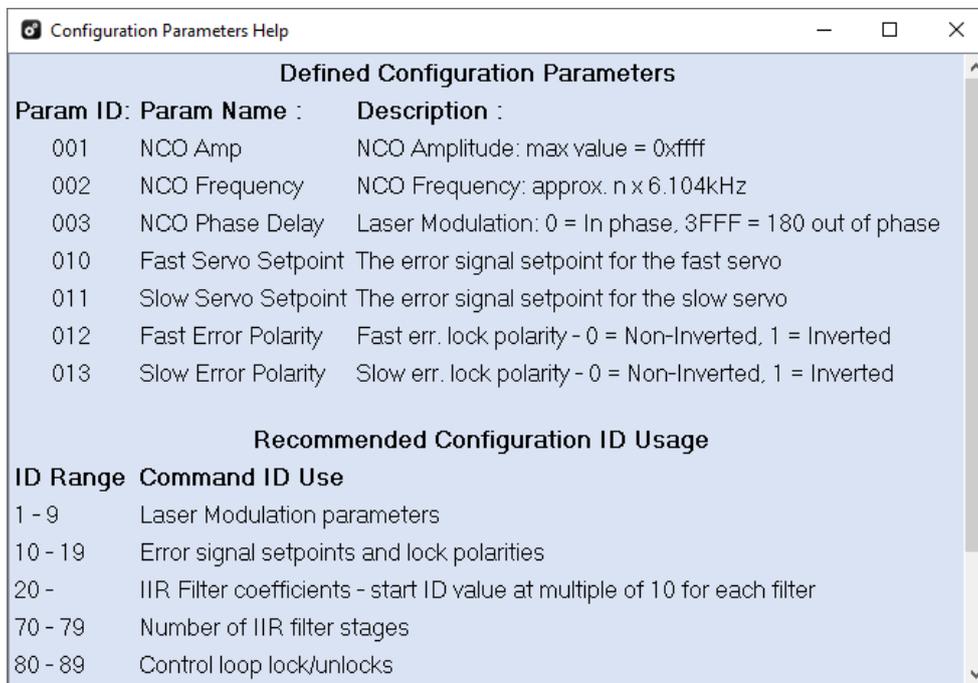


Figure 74 - 'Help' screen of the GUI interface. A description of the defined commands plus recommendations of command ID allocation is provided.

8.3.2 BACKPLANE SOFTWARE

The role of the backplane microcontroller software is to manage communications between the Front-End interface and the channel cards, provide channel card presence and lock status to the Front-End and perform any supervisory functions as needed.

The software for the ATmega324 microcontroller on the backplane is to be written as a “Bare-metal” C language application.

Each of the six channel card slots on the backplane will have a unique card presence signal routed to its own general-purpose IO pin on the microcontroller. A change of state on any of the GPIO pins can be set to trigger an interrupt, so that the microcontroller will be able to identify when a channel card has been inserted or removed and which card slot triggered this event. The microcontroller software can pass this information to the Front End.

Each channel card also has its own locked status signal, which is connected through the backplane to a unique GPIO pin on the microcontroller. Interrupts can similarly be used to identify when a particular channel has been locked or unlocked. The microcontroller software can pass this information to the Front-End or trigger a rescan of the frequency range for that channel’s laser to attempt to regain a lock.

8.3.3 CHANNEL CARD SOFTWARE

Two pieces of software have been produced to execute on one of the two ARM microprocessor cores in the Cyclone V HPS on each channel card. These are the Pre-loader (second-stage bootloader) and the application software.

When the HPS is powered up, the ROM bootloader in the HPS determines the type of boot required and the location of the Pre-loader (such as SD card or QSPI Flash) before loading the Pre-loader image into RAM and initiating the Pre-loader execution. The Pre-loader code has been generated from the open-source U-Boot software, with configuration settings based on the peripherals required by the FPGA and application software. This must be built within a Linux environment. The Pre-loader calibrates and initialises the external SDRAM and other required peripherals before copying the application software from the specified non-volatile memory source into the SDRAM and launching execution of the application software.

The channel card microprocessor application software has been written as a “bare-metal” application in the C language and makes use of API source code provided by Altera (the System-on-Chip manufacturer). The API software is split into modular blocks and provides interface functions that handle the reading and writing of the HPS and ARM processor registers.

The role of the application software is to:

- Access the FPGA programming image from the non-volatile memory device and program the FPGA.
- Manage the storage and retrieval of program data required by the FPGA to and from the EEPROM.
- Handle the reception of parameter configuration and command messages from the Backplane micro-controller or Front-End interface and pass the data or commands to the FPGA, receive status data from the FPGA and transmit FPGA status information and message responses back to the backplane micro-controller or Front-End.

8.3.4 COMMUNICATION PROTOCOL FOR PARAMETER CONFIGURATION

Configuration parameter commands from the Front-End interface will allow channel card configuration parameters (such as fast servo filter coefficients, DDS frequency and amplitude configuration and laser transmission threshold). On system power-up each channel card’s

FPGA operates from a set of pre-configured operating parameters. Parameter settings can be modified via the serial interface, but the modified settings will not be retained when the system is powered off. However, an option could be for the EEPROM on each channel card to store modified parameters and then update these in the FPGA after power up.

The serial communication protocol used to transfer parameter configuration commands between the Front-End interface or backplane micro-controller and the channel cards will convey messages in packets of ASCII printable characters. Messages will be transmitted over an asynchronous half duplex link at a data rate of either 9600bps, 19200bps, 38400bps or 115200bps with 8 data bits, one stop bit and no parity.

The Configuration Parameter message format is:

```
{<card ID>:<identifier>:<data>:<CRC>}
```

Message Field	Field Description
{	Message start sentinel, ASCII "{" (0x7B)
<card ID>	A single ASCII numeric character between 0 and 5 specifying the position of the channel card to be addressed.
<identifier>	One to three ASCII numeric characters between 0 and 9 representing a decimal command identifier in the range 0 to 127.
:	Field Separator, ASCII ":" (0x3A)
<data>	One to eight ASCII alphanumeric characters between 0-9 and A-F representing a 32-bit hexadecimal integer in the range 0 to 0xFFFFFFFF for an unsigned integer or 0x80000000 to 0x7FFFFFFF for a signed integer.
<CRC>	Eight ASCII alphanumeric characters between 0-9 and A-F representing a 32-bit hexadecimal message checksum. The CRC is calculated on the bytes in an ASCII string containing only the card ID and message identifier followed by the data field using the ISO-3309 CRC-32 algorithm on the polynomial: $x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x+1$
}	Message end sentinel, ASCII "}" (0x7D).

Table 20 - Configuration Parameter command format

Only the addressed channel card is to action and respond to any received messages. The addressed channel card microprocessor acknowledges received commands by transmitting a message in the same message format back to the Front-End host. If a correctly addressed, formatted and validated message is received by the microprocessor, the received message will be transmitted back to the Front-End. However, if a correctly addressed received message is not in the correct format or if its CRC does not match the expected CRC an error packet is returned to the front-end instead of the repeated received message.

The error packet message format is:

```
{<card ID>:<error code>:<error code>:<CRC>}
```

Message Field	Field Description
{	Message start sentinel, ASCII "{" (0x7B)
<card ID>	A single ASCII numeric character between 0 and 5 specifying the position of the channel card returning the response.
<error code>	A single ASCII character representing an error code. The following error codes are defined:

	's' : the first received character was not the expected start sentinel 'd' :the command or data field did not represent a numeric value 'c' : the received CRC was incorrect
:	Field Separator, ASCII ":" (0x3A)
<CRC>	Eight ASCII alphanumeric characters between 0-9 and A-F representing a 32-bit hexadecimal message CRC-32 checksum.
}	Message end sentinel, ASCII "}" (0x7D).

Table 21 - Configuration Parameter error response format

For example the error packet {c:c:DBB21A79} will be returned if the received message contained an incorrect CRC.

Received and validated configuration parameters will be passed to the FPGA via the HPS-FPGA bridge. Within the FPGA the command ID will be used as an index to a block of configuration data RAM where the received configuration values will be stored. This is the same block of RAM that is initialised with configuration data from the Memory Initialisation File when the FPGA is programmed.

8.3.5 CHANNEL CARD CONFIGURATION PARAMETERS

The configuration parameter commands that are planned to be implemented for each channel card are shown in Table . Light red shaded rows indicate a command rather than a configuration parameter.

Identifier	Command	Data Range	Description
01	NCO Amplitude	0 – 0xFFFF	0xFFFF: full amplitude (+/-1V)
02	NCO Frequency	0 – 0x3FFF	$f_{NCO} = \text{NCO Freq.} * f_{clk} * 2^{-15}$
03	NCO Phase Delay	0 – 0xFFFF	0: In phase, 0x3FFF: 180°
04	RF DAC Amplitude	0 – 0xFFFF	0xFFFF: full amplitude
05	RF DAC Centre Frequency	0 – $2^{32} - 1$	
10	Fast Servo Setpoint	0x80000 – 0x7FFFF	Fast error signal setpoint
11	Slow Servo Setpoint	0x80000 – 0x7FFFF	Slow error signal setpoint
12	Fast Error Polarity	0, 1	0: Non-inverted, 1: Inverted
13	Slow Error Polarity	0, 1	0: Non-inverted, 1: Inverted
20	Mixer low-pass a1	-2^{25} to $+2^{25}-1$	IIR Filter parameters for the Mixer low-pass filter
21	Mixer low-pass b0	-2^{25} to $+2^{25}-1$	
22	Mixer low-pass b1	-2^{25} to $+2^{25}-1$	
30	Fast Filter IIR coefficient a1	-2^{25} to $+2^{25}-1$	IIR filter parameters for the Fast actuator filter
31	Fast Filter IIR coefficient b0	-2^{25} to $+2^{25}-1$	
32	Fast Filter IIR coefficient b1	-2^{25} to $+2^{25}-1$	
50	Slow Filter IIR coefficient a1	-2^{25} to $+2^{25}-1$	IIR filter parameters for the Slow actuator filter.
51	Slow Filter IIR coefficient b0	-2^{25} to $+2^{25}-1$	
52	Slow Filter IIR coefficient b1	-2^{25} to $+2^{25}-1$	
70	Fast Filter MUX	0 - 15	Number of filter passes
71	Slow Filter MUX	0 - 15	Number of filter passes
80	Fast Servo Lock/Unlock	0, 1	0: Unlock, 1: Lock Valid when Test Mode = 1
81	Slow Servo Lock/Unlock	0, 1	0: Unlock, 1: Lock

			Valid when Test Mode = 1
90	Fast Servo Tune	0x80000 – 0x7FFFF	Lock tune output for the fast servo. For use in Test Mode.
91	Slow Servo Tune	0x80000 – 0x7FFFF	Lock tune output for the slow servo. For use in Test Mode.
100	Laser Transmission Threshold	0 – 0xFFFF	Threshold voltage (0V to 1V)
101	Loss of Resonance Timeout	0 - 30	Timeout in seconds after Laser Transmission falls below threshold before filters are unlocked.
110	Freq. Sweep Fast Range Min	0 – 0FFFFFF	Lower Fast sweep voltage
111	Freq. Sweep Fast Range Max	0 – 0FFFFFF	Upper Fast sweep voltage
112	Freq. Sweep Slow Range Min	0 – 0FFFFFF	Lower Slow sweep output
113	Freq. Sweep Slow Range Max	0 – 0FFFFFF	Upper Slow sweep output
114	Frequency Sweep Fast Step	0 – 0FFFFFF	Fast sweep step increment
115	Frequency Sweep Slow Step	0 – 0FFFFFF	Slow sweep step increment
120	Test Mode	0, 1	0: Normal, 1: Test mode Test Mode overrides filter locks. Also overrides Fast and Slow Tune values set by frequency sweep to use Fast and Slow Tune settings.
200	Initiate Frequency Sweep	-	Command to initiate a frequency sweep.

Table 22 - Proposed Configuration Parameter commands

8.4 FPGA FIRMWARE

8.4.1 TOP-LEVEL DESIGN

The functional design of the FPGA mixer and filter logic for each channel is shown in Figure 75 below.

The digital system mixes (multiplies) a digitized input signal from the photodiode with a demodulation waveform specific to the laser to which each channel card has been allocated to generate an error signal. The input signal is the light reflected from the reference cavity in response to modulation of that laser’s frequency. The modulation and demodulation signals are generated by twin Numerically Controlled Oscillators (NCO) with 15-bit lookup tables and an output error of +/-0.02%. To optimize the shape of the error signal, the phase of the signal used to mix down the photodiode signal must be optimized relative to the phase of the modulation. To achieve this the NCO used for the demodulation will have a phase tuning port. The two NCOs on each channel card will also have common configurable amplitude and frequency settings, which can be updated at run-time. The amplitude of the modulation can be used to optimize the error signal. An NCO clock of 200 MHz will be used, which can generate an output frequency of 10 MHz with 5 output samples per waveform quadrant. If required, the NCO clock frequency (generated by a PLL from the FPGA’s 50 MHz input clock) can be increased, allowing higher NCO output frequencies.

A low-pass filter is then used to remove the component at the mixer output at twice the modulation frequency. Finally, error signals for the fast and slow servo actuators are generated by subtracting their setpoints (nominally the equivalent of 0 V at the PD output) from the filtered signal.

The laser frequency is locked to the cavity via a fast servo and optionally also a slow servo actuator. The fast servo has been designed to support PDH loop bandwidths in excess of 1 MHz. The gain of the loop relative to the frequency is determined by the transfer function of the Infinite Impulse Response (IIR) filters, which are set by the choice of the filter coefficients. The slow servo, if required, will be able to support bandwidths in the order of tens of kHz. The loop bandwidth and tuning rate per volt of the actuator(s) on each channel card will be different, as they will be set according to the servo requirements of each laser. This will be discussed further below. Two separate digital control loops are to be used for the fast and slow loops. The error signal is split and sent through two similar systems whose parameters will be optimized separately.

First the gradient of the error signal can be switched by multiplying the signal by -1. The signal is then passed through a loop filter. It has been decided to use IIR filters to implement this. See Section 0 for details on the loop filter design.

The output of the loop filter is added to a tuning offset which could be determined from a frequency sweep of the laser assigned to each channel card. A lock/unlock parameter specifies whether the lock should be activated. The lock would be activated following a successful frequency sweep. A laser transmission photodiode input will be provided to compare the laser transmission signal strength against a user-configurable threshold value. If during normal operation the laser transmission signal strength has fallen below this signal strength threshold for greater than a configurable resonance loss time interval, the lock will be deactivated. In this case a loss of lock signal will be indicated to the backplane microcontroller which can transmit a status message to the Front-End interface. The Front-End or backplane microcontroller could then instruct the channel card to perform a new frequency scan. When a channel is unlocked its loop filter is bypassed and the DAC will follow the sum of the 'tune' parameter and the unfiltered error signal.

The loss of lock signal will also drive an LED on the front panel of each channel card, which would be illuminated when the lock is achieved.

The modulation frequency from the NCO provides phase modulation tuning words for the RF DAC around a tuneable centre frequency, which is then digitally synthesised by the DDS internal to the RF DAC. Tuning of the allocated laser frequency by ± 10 MHz is achieved by adjusting the offset frequency.

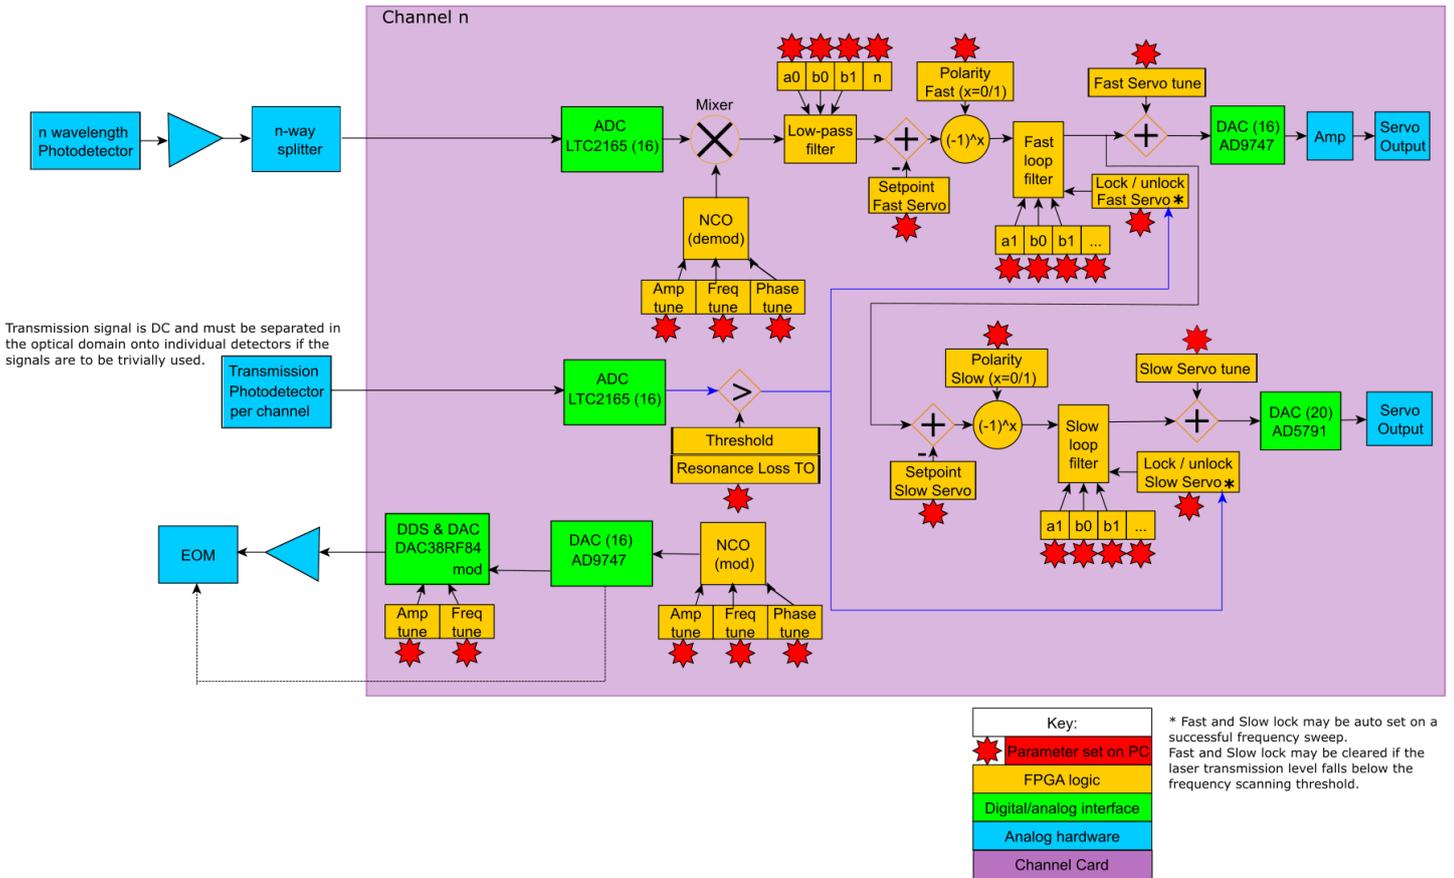


Figure 75 - Block diagram of the FPGA's digital logic (excluding frequency sweep)

8.4.2 DIGITAL FILTERS

Digital filters are required both for the low-pass on the mixer output and for the two loop filters. IIR filters are to be used instead of FIR filters as they have lower latency and require fewer FPGA resources. They also have the advantage of requiring fewer filter coefficients than FIR filters to meet the same specifications and so their transfer function can be more easily updated at run-time, as less data needs to be sent to the FPGA.

The fast and slow loop filters should be parameterized at runtime to optimize the lock of the laser to the cavity. The filter coefficients must be updateable at runtime and thus they must appear as input signals to the VHDL module. Standard implementation of first order IIR filters is well documented in the literature. The method used to design the filters here is to use the "bilinear transform method" [47].

The transfer function in the Z domain of a first order IIR filter may be written in the form:

$$\frac{Y(z)}{X(z)} = H(z) = \frac{b_0 + b_1z^{-1}}{1 + a_0z^{-1}}$$

The form of the digital logic used to perform each first order IIR filter is shown in Figure 76. Here the z^{-1} block should be read as a delay of one sample clock period. Note there is a latency of 2 sample periods in this design. This is implemented in firmware as:

$$y_n = b_0x_n + b_1x_{n-1} - a_1y_{n-1}$$

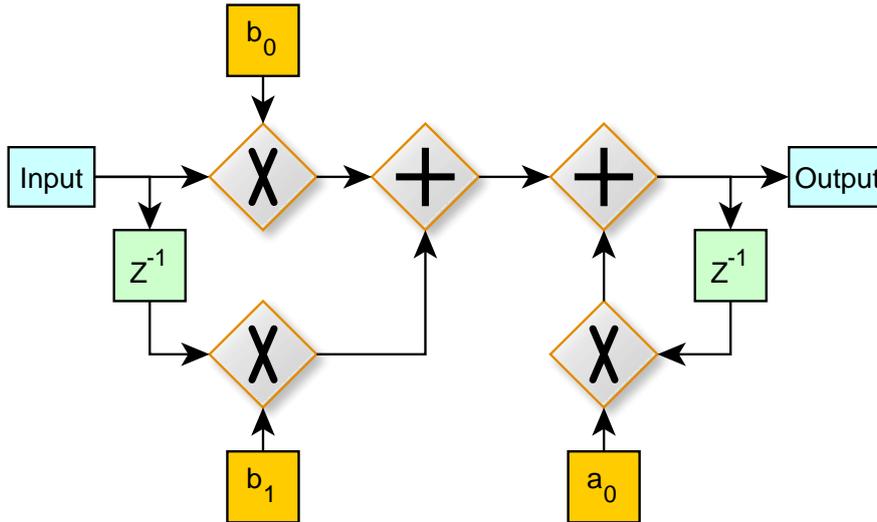


Figure 76 - The first order IIR filter design to be implemented on the FPGA

Individual IIR filters will be combined either in parallel or in series to create a system with a more complex transfer function. A series combination of five first order IIR filters is to be used in the configuration shown in Figure 77.

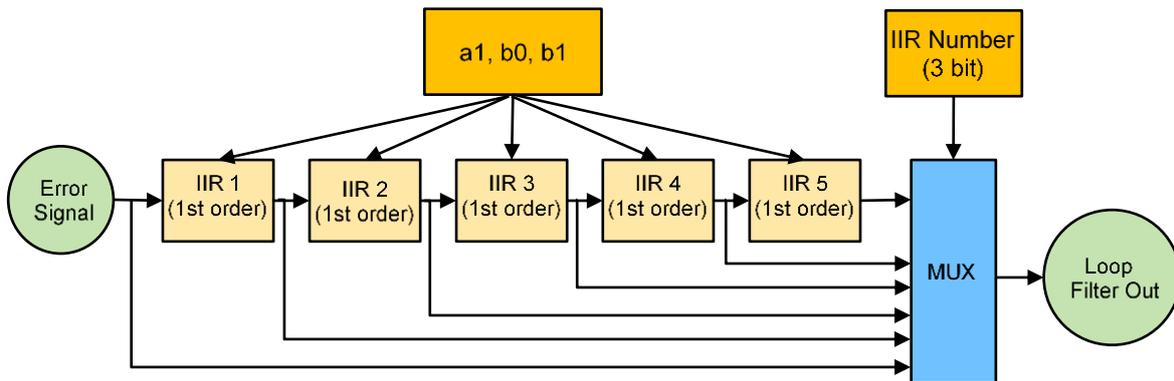


Figure 77 - Block diagram of the digital filter to be used for both fast and slow feedback. A MUX allows to take the output of any of the IIR filter stages to reduce the overall latency if a less complex filter is required.

To obtain the overall filter transfer function, the phase response of the individual filters should be summed while the amplitude response is multiplied. Note if the amplitude response is obtained in dB then the individual amplitude responses should be summed.

The IIR filters can be used to implement a PI controller. The transfer function of a PI controller in the Z domain may be expressed in the form:

$$\frac{Y(z)}{X(z)} = H(z) = K_p + \frac{K_I T}{2} \left(\frac{z + 1}{z - 1} \right)$$

Manipulating the two transfer functions above, the IIR filter coefficients can be shown to be:

$$a_1 = -1, b_0 = K_p + \frac{K_I T}{2}, \quad b_0 = -K_p + \frac{K_I T}{2}$$

The actual loop filter transfer function to be used in the system will be optimized when the device is ready to be tested. The slow servo system will have a pole at the cavity linewidth. Poles at low frequency will be used to create a high gain and zeros at a higher frequency will keep the system stable.

The latency in the implemented firmware from ADC sampling of the PDH signal through to the fast servo DAC output is 190 ns. This has been calculated by summing the clock delays through each FPGA module, with three passes through the low pass filter and one pass only through the fast servo PI filter. There will be a 180° phase shift at a frequency of $1/(2 \times \text{total latency}) = 2.63 \text{ MHz}$ [48], which limits the fast servo bandwidth to around 2 MHz

The linewidth of 1 MHz to which the 461 nm laser is required to be locked (R-03 in [20]) is close to the laser's free-running linewidth, and so the servo will mainly need to provide frequency drift control, likely to be at a bandwidth of <1 kHz. The same servo bandwidth will be needed for the 679 nm and 707 nm lasers, which will also need only frequency drift control. The servo bandwidth of the 813 nm laser will be in the order of hundreds of kHz. The specified linewidths of the 689 nm laser (R-03 in [20]) of 1 kHz will require low MHz servo bandwidths.

8.5 VERIFICATION OF PDH LOCK

To validate the use of a digital fast PDH servo lock scheme for laser stabilisation, control electronics and software have been implemented to demonstrate servo locking of a single laser frequency through the optical cavity. The control electronics' FPGA firmware has been executed on a PC under a simulation environment to model the operation of locking scheme components, in addition to the FPGA firmware running on the control electronics. The sections below describe, for both the simulation and electronics environments, the set-up of the demonstration followed by presentation and discussion of the results to validate the design of the digital servo control loop.

8.6 LOCKING SCHEME – FPGA SIMULATION

8.6.1 DEMONSTRATION SET-UP

The FPGA simulation executes a subset of the FPGA firmware developed for the MW-CCU channel cards. Instead of running on the Cyclone V FPGA, the firmware is interpreted by Mentor Graphics' *ModelSim* simulation environment executing on a PC. 'Testbench' FPGA firmware has been created to simulate the FPGA device reset, clock inputs, photodetector input and set configuration parameters. The *ModelSim* software package used is the *ModelSim* Starter Edition 2020.1, revision 2020.02

The FPGA firmware executed in the simulation covers the NCO generation of the demodulation waveform, the mixing of this waveform with the digitised photodetector signal and the low pass filtering of these mixed signals (these modules can be seen in Figure). The second NCO, normally used for generating the modulation signal, generates a simulated digitised photodetector input.

The simulation sets the demodulation NCO for an output frequency of 10 MHz, at full amplitude and no phase offset. The digitised photodetector input NCO is set for an output frequency of 10 MHz + frequency offset, where this offset could range from $\pm 25 \text{ kHz}$ to $\pm 1200 \text{ kHz}$ in steps of 25 kHz. Here, the output amplitude is set to 0.079% of full scale, to represent the amplitude of the ADC output when a $\pm 2.5 \text{ V}$ photodetector input is applied to the control electronics.

A simulation was run for each photodetector input frequency over a chosen range and a resultant waveform display was generated to show the demodulation signal, digitised photodetector signal, mixer output, output of each of the three cascaded stages of the low pass filter, and a 16-bit resolution copy of the final stage low pass filter output. The maximum amplitude of the low pass filter output was then logged for each photodetector input frequency, to produce a graph of filter output amplitudes against photodetector to demodulation frequency offsets. This enabled the 3 dB cut-off frequency of the low pass filter to be determined.

8.6.2 RESULTS DISCUSSION

The ModelSim-generated waveform diagram of an FPGA simulation with 10.4 MHz digitised photodetector input is shown in Figure below.

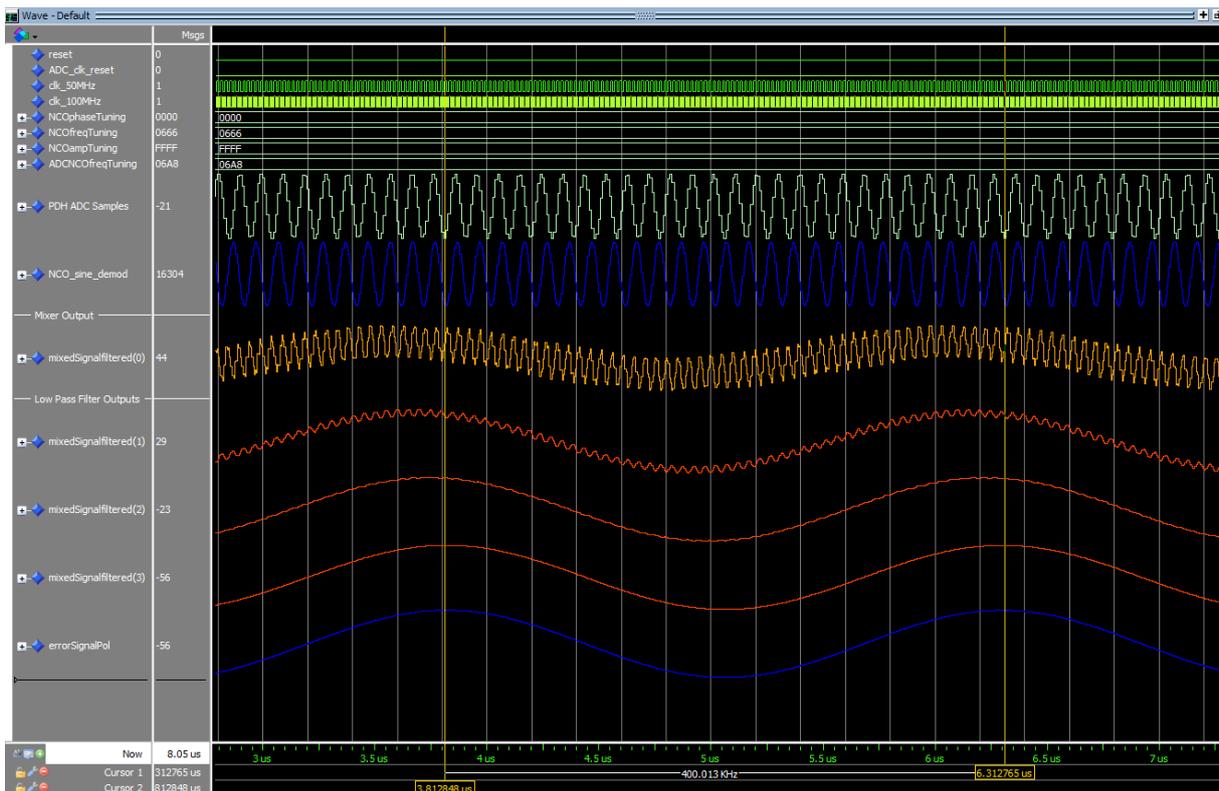


Figure 78 - ModelSim waveform diagram for 10.4 MHz photodetector input

It can be seen in Figure 78 that the resolution of the digitised photodetector input (in green) is much lower than that of the demodulation NCO output (in blue). The demodulation NCO is clocked at 200 MHz, while the photodetector NCO is driven at only 50 MHz, representing the input clock to the LTC2165 ADC.

The low pass filter output (bottom waveform in green) has a frequency equal to the frequency difference between the input photodetector signal and the demodulation waveform, in this case 400 kHz. The three waveforms in red are the outputs of the first, second and third passes through the low pass filter. In the subsequent modules of the FPGA firmware, the error signal is input to the digital servo, which will act to reduce the frequency difference between the demodulation signal and the photodetector input.

As the input photodetector frequency is increased, the amplitude of the low pass filter output can be seen to decrease (see Figure 79 below). Because the low pass filter has been implemented as a first order digital IIR filter, its frequency response can be adjusted by

changing the a_1 , b_1 and b_0 coefficient values. The lowpass filter can then be calibrated to have a 1 MHz bandwidth via the selection of suitable filter coefficients.

The simulations that produced the waveform displays shown in Figures 78 & 79 used the following coefficients: $a_1 = 0.95310$, $b_1 = 0.03779$, $b_0 = 0.03779$. Figure shows the frequency response of the low pass filter output waveform amplitudes resulting from each input photodetector frequency when the NCO demodulation frequency was set to 10 MHz.

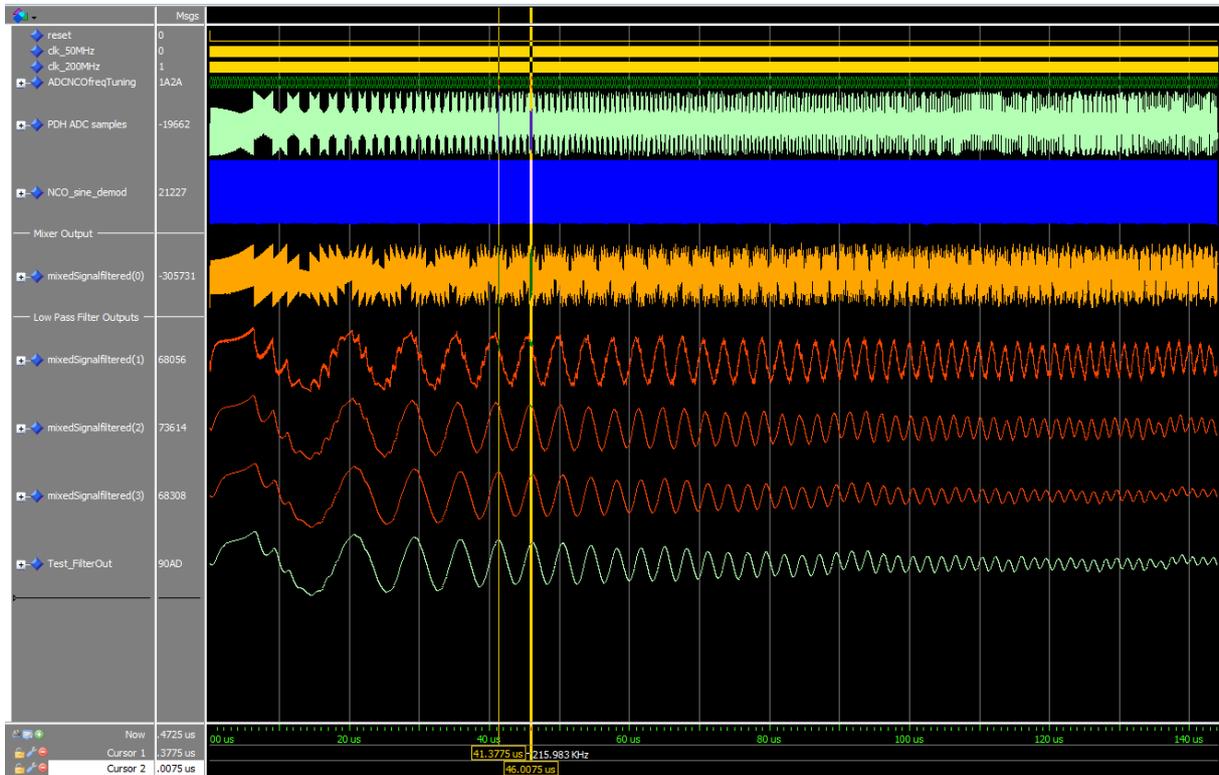


Figure 79 - ModelSim waveform diagram with a sweep of the photodetector input frequency

The filter coefficients chosen here result in a smooth roll-off of the filter output amplitude with increasing frequency. Figure 80 shows a low path filter bandwidth of approximately 1.275 MHz, but the filter coefficients could be optimised for a different cut-off frequency and/or filter gain if desired.

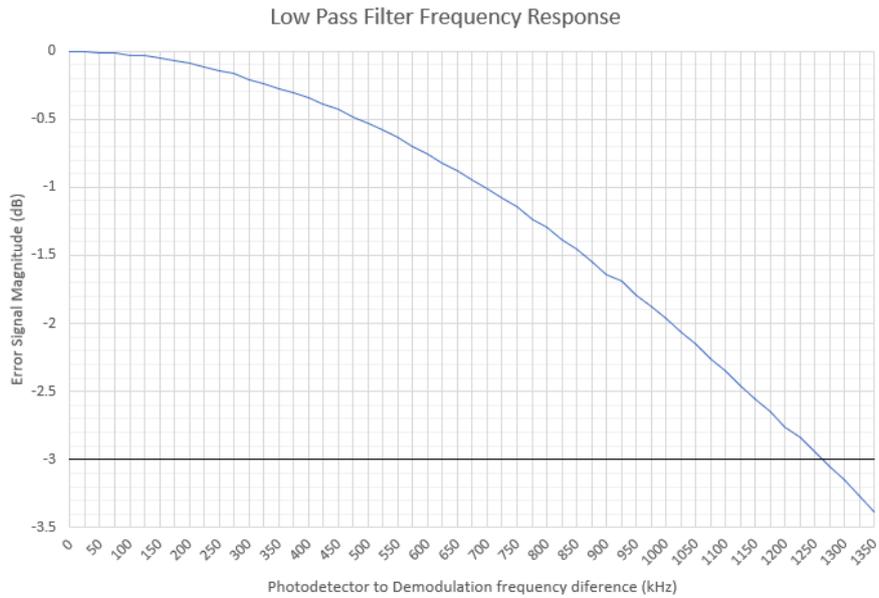


Figure 80 - Frequency Response of Low Pass Filter

8.7 LOCKING SCHEME – ELECTRONICS IMPLEMENTATION

8.7.1 DEMONSTRATION SET-UP

The electronics hardware used for the demonstration is the NPL-designed PCB. This was chosen because the MW-CCU channel card design has been based on existing control electronics, with the exception of the RF DAC for offset sideband modulation, which is not required here. The FPGA firmware from boards developed within parallel NPL projects has been modified to support the NCO output frequency required for CCU modulation/demodulation and the higher servo bandwidth requirements.

The FPGA firmware executed in the demonstration covers the NCO generation of the demodulation waveform, the ADC sampling of the photodetector input signal, the mixing of the demodulation waveform with the digitised photodetector signal, the low pass filtering of these mixed signals and the DAC output of the low pass filter.

A function generator has been used to produce a +/-2.5V sinewave that is applied to the photodetector signal ADC (LTC2165) input on the PCB. As with the FPGA simulation, the demodulation NCO has been set for a demodulation frequency of 10 MHz at full amplitude and with no phase offset. The upper 16 bits of the low pass filter's 20-bit output is sent to one of the 16-bit DAC (AD9747) channels and the output of the DAC is monitored by an oscilloscope.

The function generator has been set to output the photodetector signal with a frequency of 10 MHz + frequency offset, where this offset can range from 25 kHz to 1200 kHz in steps of 25 kHz.

The low pass filter output amplitude and frequency is to be logged for each function generator frequency, to produce a graph of filter output amplitudes against photodetector to demodulation frequency offsets. This enabled the 3 dB cut-off frequency of the low pass filter, and hence, its bandwidth to be determined.

The equipment used for the demonstration is:

- Control electronics PCB (Board ID 110902), with DAC output offset set to 0V (see Figure).

- Front End Interface program, version 1.02.
- FeelTech FY6900-60M DDS Function/Arbitrary Waveform Generator, firmware V1.3.1 (S/N 902007990)
- Tektronix MDO3024 Mixed Domain Oscilloscope, 2.5 GSPS maximum (S/N C033170)

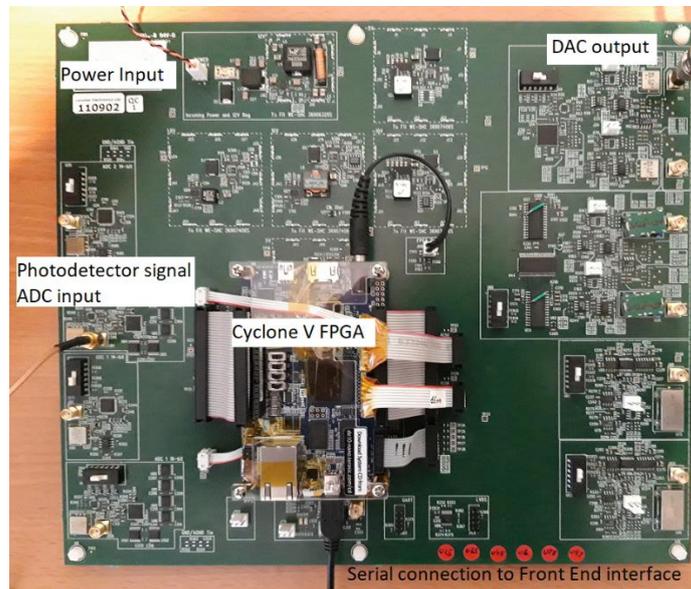


Figure 81 - Electronics PCB used for MW-CCU locking demonstration

8.7.2 RESULTS DISCUSSION

Inspection of the mixed demodulation and 50 MHz-sampled photodetector signals revealed a noisy waveform that was unlikely to produce filtered successfully. Instead, a 100 MHz clock signal was applied to the LTC2165 ADC for sampling the photodetector signal.

The oscilloscope trace of a 10.6 MHz sinewave from the function generator mixed with the 10 MHz demodulation signal is shown in Figure 82. Waveform 1 (yellow) is the 10.6 MHz sinewave output from the function generator. Waveform 2 (blue) is the mixed signal. The large oscillation in this trace is at a frequency of 600 kHz, being the frequency difference between the 10.6 MHz function generator signal and the 10 MHz demodulation waveform. The low pass filter will work to remove the smaller oscillations (at 20.6 MHz = 10 MHz demodulation signal + 10.6 MHz function generator sinewave), leaving only the 600 kHz waveform.

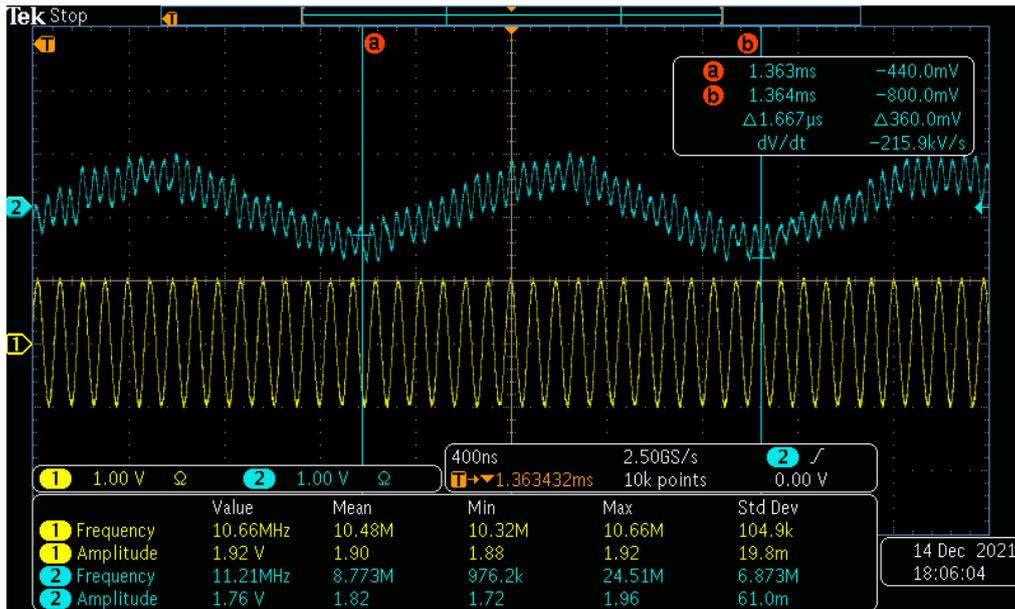


Figure 82 - Mixed 10 MHz demodulation and 10.6 MHz photodetector waveform

8.8 EVALUATION OF THE SERVO USING A 1064 NM DFB LASER

In this section, we describe tests undertaken to evaluate the servo using an Eagleyard Photonics free-space DFB laser (part EYP-DFB-1064-00080-1500-TOC03-0002) which has a specified FWHM linewidth of 2 MHz. The laser was driven using an NPL version of a published low-noise current source [49]; the output spectrum was obtained by beat frequency comparison (Figure 83) against a sub-Hz linewidth reference laser. This reference laser is locked to a 48-cm cavity and was part of a recent European comparison of clock stabilities [50]; this reference gives details of the NPL system in the “methods” section.

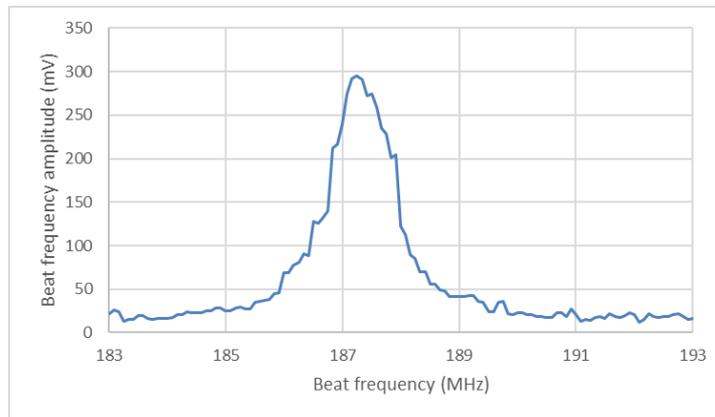


Figure 83: Beat frequency between the DFB laser used in this TN and a sub-Hz linewidth reference laser. The beat linewidth is ~1.2 MHz FWHM.

The electronics built and tested in this work-package are one of a pair required for the later “OSRC testbed” project⁵. This project will develop a dual axis cubic cavity; one arm is specified for use at 1397 nm and the second for use at 1064 nm. The frequency stability of the 1064-nm bore will be demonstrated by locking a YAG laser. A second electronics board with a lower bandwidth and ~2 MHz modulation frequency will be developed for this application. The electronics in this TN will later be used for locking a Toptica DL-PRO 1397 nm laser. The DL-PRO specified linewidth is <300 kHz at 5 μs⁶, less than the 1.2 MHz DFB linewidth at 1064 nm

5 ESA contract 4000122879/18/NL/PS/zk

6 <https://www.toptica.com/products/tunable-diode-lasers/ecdl-dfb-lasers/dl-pro>

and so the eventual servo bandwidth demands should not be as stringent. The electronics developed here use a modulation frequency of ~ 10 MHz, a low pass filter bandwidth of 1.25 MHz (using the same filter coefficients as in the FPGA simulation (Figure 80) and have an expected servo bandwidth of ~ 2.0 MHz. For these tests, a cubic cavity was used that was developed within an earlier ESA activity.

The performance of these servo electronics was monitored by measuring the cavity transmitted signal with the servo active. The servo IIR filter coefficients used were: $a_1 = 1$, $b_1 = -6.0797 \times 10^{-5}$, $b_0 = 6.1511 \times 10^{-5}$. These generated a proportional gain of 6.115×10^{-5} and weak integral gain of 71.526 and, with these settings, the laser frequency was controlled near the cavity resonance for ~ 20 ms (Figure 84). Additional optimisation of both the high frequency and integral gains would be required for these electronics to control a laser frequency at the sub-Hz level, as required for the 1397 nm ECDL in “OSRC testbed”.

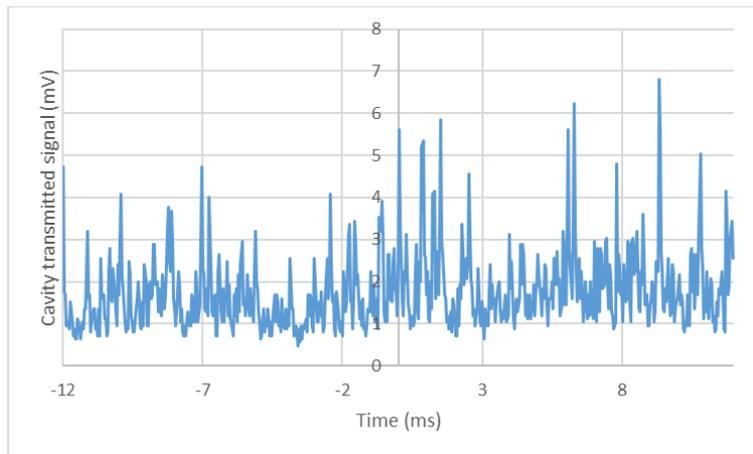


Figure 84: Signal transmitted through a 5-cm cubic cavity using the newly developed digital servo with ~ 2 MHz bandwidth. Further optimisation will be required before final deployment and this will be undertaken with a 1397 nm ECDL in the later “OSRC testbed” project

In summary for the sections 8, 9 & 10, we have provided a detailed design of updated control electronics to include a fully digital fast PDH servo locking scheme for each of the laser frequencies to be stabilised in the MW-CCU. The electronics hardware has been designed as a modular system executing a configurable software solution that provides the flexibility to control and stabilise up to six lasers, each with their own particular modulation frequency and servo bandwidth requirements.

To support the digital PDH servo design proposal, a digital servo lock has been demonstrated using both a simulation environment and the PDH control electronics hardware from the NPL boards with updated FPGA firmware to provide the higher modulation frequency and servo bandwidth required for CCU clock laser stabilisation.

9 MANUFACTURING PLAN FOR AN ENHANCED CCU

9.1 INTRODUCTION

In this concluding section we develop the implementation plan for the manufacture, assembly, design and test (MAIT) of the CCU according to the updated designs articulated within this activity. The CCU MAIT will be undertaken in a relevant follow-on activity. We outline plans for improved performance of the NPL dual-axis cubic cavity and review our plans for:

- A determination of the current manufacturing tolerances using measurements currently underway on a NPL CMM followed by modelling to understand whether the target

tolerances are consistent with the observed cavity sensitivity to acceleration exceeding $10^{-10}/g$ for recent builds in some directions.

- An improved procedure for the manufacture of the cavity mounting frame with a process to align and centre the cavity at the low tens of microns level rather than the current uncertainty of 100 μm . We think it is likely that the range of sensitivities to acceleration in recent builds can be explained by manufacturing tolerances. We also consider whether the range of vibration sensitivities could be explained by the range of torques used to secure the cubes within the mounting frame.
- A brief review of the changes required to upgrade the clock axis to a high finesse at 1397 nm using crystalline mirrors
- A review of the potential advantages of orientating the ULE dual axis cube so that both axes are perpendicular to the ULE growth axis
- A future CCU dual axis cavity will incorporate designs and outputs from parallel ESA projects (including OSRC, testbed and HSL2) into a single device.

9.2 COMPACT CUBE OPTIONS

In Section 6 we explored different options for a smaller cubic cavity; both 3-cm and 4-cm cube options were considered. For example, Table 13 includes options for a 4-cm cube and $\frac{1}{2}$ " diameter fused silica mirrors. The manufacturability of some of these options needs to be explored further with suppliers to down-select to a more compact cube option. The design needs to provide both a minimised reduction in the effective zero CTE temperature when using fused silica mirrors and a minimised sensitivity to acceleration. For NGGM rather than optical clock applications which require short-term frequency stabilities in the parts in 10^{16} region, the use of ULE rather than fused silica mirrors could allow us to explore further small cavity configurations.

9.3 IMPROVED CAVITY ALIGNMENT PROCEDURE

In this section, we outline some recent improvements to our cavity manufacturing and alignment process. For future cavity builds, optical contacting will be carried out by attaching an alignment jig (see Figure 85) to the surface of the cube. This jig is designed to fit a 5 cm cube and provides a 25.4 mm centrally aligned hole into which the mirror can be placed. Gentle, even force is applied by hand to press the mirrors onto the cube and create the optical contact. Contact regions are manually inspected for marks or blemishes, and mirrors re-cleaned and re-contacted where necessary.

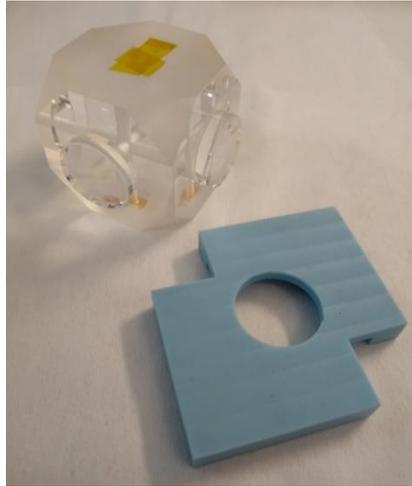


Figure 85 – Cubic cavity with optically contacted mirrors (free bores covered with Kapton tape), alongside the alignment jig

Once sufficiently high finesse values are obtained, the cavity will be mounted within a gold-plated aluminium frame. We note here that thermal modelling by STI as part of OSRC indicate that gold coating of the frame and heatshields provides an improved thermal performance. Alignment will be carried out using micrometer jigs attached to the cubic frame via four M4 bolts. Each jig consists of a steel plate with a pair of Mitutoyo 148-133 micrometer heads (10 μm graduations, accurate to $\pm 2 \mu\text{m}$) mounted within them. The design position of the cube faces from the edges of the frame is 17.5 mm, and the micrometer heads can then be set to $\pm 5 \mu\text{m}$ of this using a bespoke calibration block (see Figure 86).

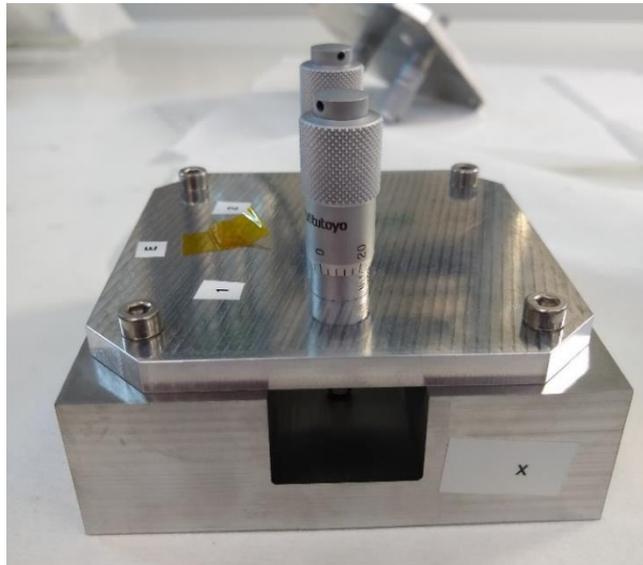


Figure 86 – Initial calibration of the micrometer jigs to the expected cube distance prior to being fixed onto the cubic frame

Once set to a length corresponding to the expected position of a cube face relative to the frame edge, three of these micrometer jigs (attached to three faces adjacent to the same corner of the frame) are used to position the cube accurately.

With the frame resting on its corner the cube was rested on the micrometer heads, with the six points of contact being sufficient to hold the cube in the correct position and orientation. The four spherical supports are wound down onto the truncated vertices so that the PEEK spheres and the cube corners just come into contact, before being torqued to 0.3 Nm. Assuming a bolt

coefficient for standard, non-lubricated steel fasteners of $K = 0.2$ [51], this corresponds to a contact force, P , between the PEEK spheres and the cubic cavity at each support of

$$P = \frac{T}{K \cdot d} = 250 \text{ N}$$

where T is the bolt torque and d is the bolt diameter. This contact force level is chosen based on previous cubic cavity builds [52]. Once the torques are even, the micrometer jigs can be removed.

At this stage the ULE annuli are optically contacted onto the back faces of the mirror substrates. In the current design the cavity is not able to fit inside the frame after the ULE annuli are contacted, so the cube (with mirrors) must be mounted in place first. This enables the cubic frame to be smaller and make the overall vacuum chamber and LSU system more compact. The cube position is re-measured to check that the contacting process has not caused any misalignment.

9.4 PLANNED CMM MEASUREMENTS BY NPL DIMENSIONAL TEAM

Measurements of a recently constructed cavity mounting frame by the NPL dimensional team are underway and will be reviewed within another ESA project when completed. The results are expected to contribute to understanding whether manufacturing tolerances, currently $100 \mu\text{m}$, could be contributing to the observed range of acceleration sensitivities observed [40].

9.5 SELECTION OF ULE ORIENTATION

As discussed in [40], the original CCU was manufactured with one axis parallel to the growth direction. This means that the clock and auxiliary axes have different thermal properties; for example, the temperature at which the CTE is zero is different for the two directions. For this application, a future CCU might have a better performance if both directions are perpendicular to the growth axis. This would mean that the temperature at which the CTE is zero would be the same for both directions. The thermal expansion in the two directions would also be similar making feedforward drift control more straightforward. We might also expect the isothermal drift to be similar for the two directions, but this needs to be investigated further.

9.6 UPGRADING THE CCU CLOCK AXIS

Requirements for upgrading the CCU clock axis to a high finesse cavity with the option of using crystalline mirror coatings at 1397 nm were discussed earlier. Further considerations that will determine the final optics configuration will include:

- Whether to work with ThorLabs/ CMS to develop a crystalline mirror coating [13] for the 674-nm sub-harmonic at 1348 nm . This will be required if the CCU needs to target a Sr^+ ion [4] rather than a Sr lattice clock [43]. An ideal outcome would be a crystalline mirror coating to cover both 1348 nm and 1397 nm .
- Whether to implement fibre path length stabilisation [46]; this will need to be decided on the basis of the environmental thermal noise and length of fibre between the laser, cavity and clock physics packages. For the most demanding applications, implementation of path length stabilisation is almost certainly required.
- Finally, we need to understand whether active RAM control [45] is required. For the most demanding applications, this will almost certainly be a requirement.

10 IMPLEMENTATION PLAN

This section contains further details of the planned implementation outlined below. This section concludes with a roadmap in section 10.2 and budgetary costing in section 10.3.

TASK 1: CUBIC CAVITY & MOUNTING REFINEMENT STUDIES & DEVELOPMENT

- **Choice of DACC axes relative to ULE growth axis**
This activity will examine the difference in ULE cavity isothermal drift rates for the two orthogonal cavity axes for the situations where one axis is along the ULE boule growth axis and one perpendicular to that axis – i.e. radially from the growth axis, against the difference for both axes along orthogonal radial directions. Currently, an assumption is that the drift rates would be similar in the latter case, but this needs to be confirmed, and the level of agreement determined.
- **Cube machining tolerances (linear, squareness, angle), and effect on insensitivity**
Complete a programme of cube dimension, squareness & angle by measurements made on NPL's high accuracy co-ordinate measuring machine CMM capable of few microns uncertainty in 3D
- **Cavity mirror alignment optimisation techniques**
Taking account of the CMM measurements, refinement of the mirror alignment procedure.
- **Mounting frame machining tolerance (ditto) and effect on asymmetry & insensitivity**
Further CMM measurements on the fidelity of 3D dimension, squareness, angle of the mounting frame and the analysis of their effect of the resulting asymmetry on the force and vibration insensitivity.
- **Frame-to-cube vertex force insensitive loading studies and refinement**
Tests of different loading levels on the vertex mounting screws and their effect on vibration insensitivity
- **Cube size reduction and effect on asymmetry**
Following earlier FEA on cube size reduction below 5 cm, selection of relevant reduced size and machining and preliminary testing of vibration insensitivity
- **Reduction of asymmetry due to mirror attachments, investigation of dummy mirrors on 3rd axis**
Apply dummy mirror substrates (and annuli if other axes have these) to 5 cm cube to reduce asymmetry and investigate whether this leads to improved vibration insensitivity

TASK 2: VACUUM HOUSING REFINEMENT

- **Vacuum housing size, materials, sealing, bake-out parameter refinement**
Mindful of data gathered from existing cubic cavity modelling analyses and developments, a cavity housing will be built that takes account of this prior data in respect of the materials, sealing arrangements and bake-out procedures to optimise the achievable UHV low pressure.
- **Ion pump / getter combinations & resulting vacuum levels achievable**
Consideration will be given to the question of whether the cavity chamber can be sealed off with just getter pumping operational, or whether the ion pump is absolutely necessary. This has implications on the frequency sensitivity of the complete cavity & housing structure to vibration where the ion pump hangs off the side of the chamber

housing. Additionally, the option of a vent to space, once the system would be deployed in orbit, as was done in the LISA Pathfinder demonstrator, will be examined.

- **Cube mounting frame to chamber fixing for minimum vibration / shock insensitivity. central ring mount design/alternative for minimum low frequency resonances.**
The finite element analysis and review conducted within this contract has given rise to a perceived need for improved fixing of the cube and mounting frame to the vacuum chamber, in order to cope with the tendency of the cube/mounting frame to demonstrate increased vibration (“wobble”) due to its fixing only to the chamber base. Two potential solutions are suggested. The first would involve the inclusion of downward pressure bearing down onto the top surface of the mounting frame by means of PEEK stubs fixed into the inner roof of the vacuum chamber. The second would involve a square that allows mounting ring providing mid-plane support similar to the OSRC case. The slight difficulty in this latter approach is that the mid-plane is also the dual axis cavity plane, necessitating a ring design with light throughput on the two axes and a small increase in vacuum chamber size to accommodate the ring.

TASK 3: OPTICS REFINEMENT FOR COMPLETE CLOCK AND AUXILIARY INSENSITIVE SYSTEM

- **Dual axis with very high finesse coating on axis for clock stabilisation**
In order to upgrade the original CCU clock reference axis to provide the primary stabilisation of the clock laser, the clock axis coating finesse has to be increased from a few thousand to ~ 200,000, ideally by means of a low thermal noise crystalline mirror coating. Due to the current non-availability of crystalline mirror coatings at 698 nm, this will likely necessitate the clock laser operating at the half-wavelength of the clock frequency (e.g. 1397 nm for the 698 nm Sr lattice clock frequency) with subsequent doubling to 698 nm. In order to cater also for a Sr⁺ ion clock system with a clock laser at 674 nm, the optimum design wavelength for the crystalline coating would be 1371 nm, halfway between the Sr neutral and Sr⁺ ion clock half-wavelengths.
- **Auxiliary axis coating capable of dealing with both Sr and Sr⁺ ion systems**
Given the desirability of a CCU design that can cater for both Sr and Sr⁺ clock systems, the auxiliary axis coating bandwidth will need to cater for the clear-out lasers at 1033 nm and 1092 nm. This is considered viable given the lower finesse levels needed.
- **Integration of universal optics arm (possible delta variant needed?)**
The need for a variant design for the universal EQM optics arm will need to be considered, given the integration of multiple laser inputs into the optics arm assemblies via collimator/fibre coupler arrangements.

TASK 4: OPTO-ELECTRONICS & OPTICS

- **Integration of multiple EOMs for selective PDH modulation of separate wavelengths into fibre/coupler feeds**
- **Integration of optics arrangements for RAM, phase noise stabilisation**

TASK 5: SERVO-ELECTRONICS DESIGN REFINEMENT

- **High frequency modulation and bandwidth for PDH clock laser stabilisation**
- **Digital servo control and system control throughout**
- **Autolock and recovery software**

TASK 6: INTEGRATION AND FULL UNIT PERFORMANCE TEST & ENVIRONMENTAL TEST

- Stability, stability floor, long term stability / drift
- Environmental testing (vibration/shock, thermal vacuum/ radiation) to ESA requirements

TASK 7: ADVANCE TO TRL7 VIA ENVIRONMENTAL TESTING PROGRAMME FOR FREE FLYER

- Preliminary considerations of space qualification requirements of CCU system and sub-components
- Achievement of Engineering model of complete system for possible launch on GSTP 6.3 In-orbit technology demonstrator as a prelude to full EQM/FM

10.1 CURRENT STATUS WITHIN DIFFERENT ESA PROJECTS

In the medium term, it is anticipated that the out-turn to a 3-year follow-on project to the current MW-CCU design contract would be an EM MW-CCU model tested to TRL7 with acceleration insensitivity in the low parts in $10^{11}/g$ along all three axes for the dual axis cube. In this respect, we plan on similar acceleration performance to the original single axis cube. Whilst modelling indicates that this is possible, the practicality of manufacturing to tight tolerances along an increased number of parameters versus the single axis cube has proved challenging [40]. With this achievement, it can be considered for an in-orbit demonstrator / in-orbit validation activity (IOD/IOV) on board a free flyer as a part or full payload under the GSTP 6.3 programme.

Input to a future dual axis CCU follow-on from existing projects is expected to include:

- A review of the OSRC design, in particular consideration of a mid-plane ring mounting arrangement that emulated the OSRC ring design that has demonstrated to TRL6
- The outcome of the delta design review in Testbed, in particular the availability, specifications, and planned use of dampers.
- The design of a EQM optics arm that is starting as CCN3 in the HSL2 project. This design could be transferred with potentially minor modifications in view of the change of wavelength (1064 nm to 1397 nm) for the clock arm but more modifications will be required for the other clocks wavelengths to allow these to be multiplexed into the cavity.

10.2 DEVELOPMENT ROADMAP

The planned steps within a future development is summarised in Figure 87.

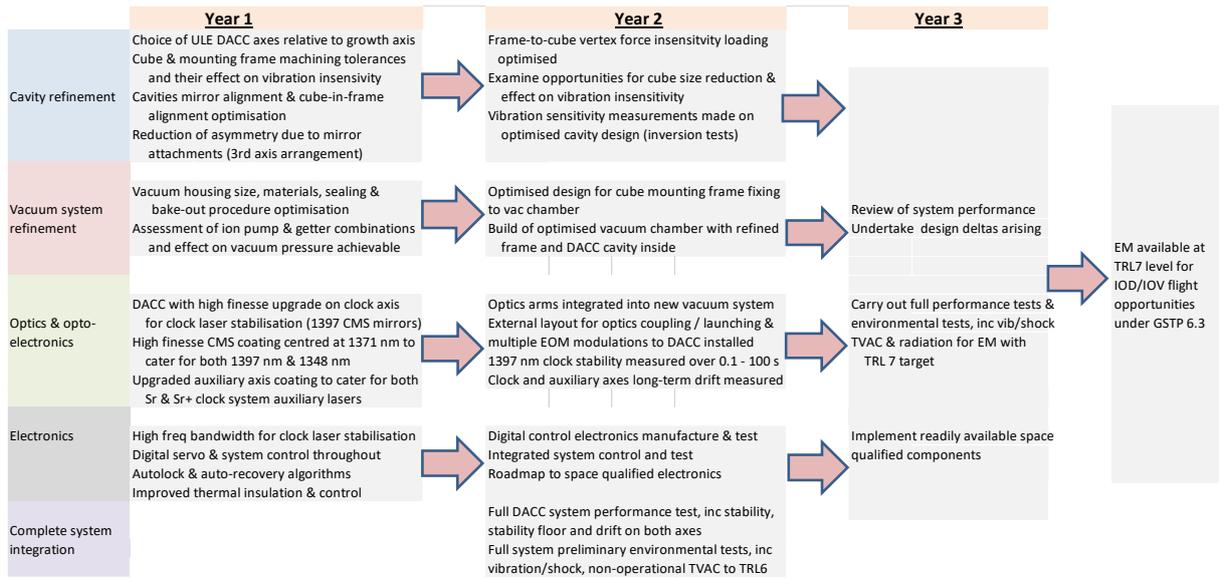


Figure 87: Development roadmap for a future dual axis EM cubic cavity

10.3 BUDGET ESTIMATE FOR PROPOSED ACTIVITIES

The anticipated cost for the proposed activities on a 3-year timescale to reach TRL7 is estimated to be 800 k€. This covers only estimated NPL costs. Large space integrator LSI involvement is assumed, but where additional costs for the LSI are not itemised here.

In the medium term, it is anticipated that the out-turn to a 3-year follow-on project to the current MW-CCU design contract would be an EM MW-CCU model tested to TRL7. With this achievement, it can be considered for an in-orbit demonstrator / in-orbit validation activity (IOD/IOV) on board a free flyer as a part or full payload under the GSTP 6.3 programme.

Subsequent to this follow-on activity, in the longer term it is envisaged that the MW-CCU system could be implemented in a future Space Optical Clock mission, and in the Earth Observation NGGM and LISA Science missions and other Fundamental Science opportunities.

REFERENCES

- [1] S. A. Webster and P. Gill, “Force-insensitive optical cavity,” *Optics Letters*, vol. 36, pp. 3572-4, 2011.
- [2] R. Hobson, W. Bowden, A. Vianello, I. R. Hill and P. Gill, “Cavity-enhanced non-destructive detection of atoms for an optical lattice clock,” *Optics Express*, vol. 27, p. 37099, 2019.
- [3] R. Hobson, W. Bowden, A. Vianello, I. R. Hill and P. Gill, “Midinfrared magneto-optical trap of metastable strontium for an optical lattice clock,” *Phys Rev A*, vol. 101, p. 013420, 2020.
- [4] G. P. Barwood, G. Huang, H. A. Klein, L. A. M. Johnson, S. A. King, H. S. Margolis, K. Szymaniec and P. Gill, “Agreement between two 88Sr+ optical clocks to 4 parts in 10¹⁷,” *Phys Rev A*, vol. 89, p. 050501(R), 2014.
- [5] G. P. Barwood, P. Gill, G. Huang and H. A. Klein, “Automatic laser control for a 88Sr+ optical frequency standard,” *Measurement Science & Technology*, vol. 23, p. 055201, 2012.
- [6] T. Fordell, T. Lindvall, P. Dubé, A. A. Madej, A. E. Wallin and M. Merimaa, “Broadband, unpolarized repumping and clearout light sources for Sr+ single-ion clocks,” *Optics Letters*, vol. 40, pp. 1822-5, 2015.

- [7] A. Nevsky, S. Alighanbari, Q.-F. Chen, I. Ernsting, S. Vasilyev, S. Schiller, G. P. Barwood, P. Gill, N. Poli and G. Tino, “Robust frequency stabilization of multiple spectroscopy lasers with large and tunable offset frequencies,” *Optics Letters*, vol. 38, pp. 4903-6, 2013.
- [8] T. Legero, T. Kessler and U. Sterr, “Tuning the thermal expansion properties of optical reference cavities with fused silica mirrors,” *J. Opt. Soc. Am. B*, vol. 27, pp. 914-9, 2010.
- [9] J. Sanjuan, K. Abich, M. Gohkle, A. Resch, T. Schuldt, T. Wegehaupt, G. P. Barwood, P. Gill and C. Braxmaier, “Long-term stable optical cavity for special relativity tests in space,” *Optics Express*, vol. 27, p. 36206, 2019.
- [10] S. Häfner, S. Falke, C. Grebing, S. Vogt, T. Legero, M. Merimaa, C. Lisdat and U. Uwe Sterr, “ 8×10^{-17} fractional laser frequency instability with a long room-temperature cavity,” *Opt Lett*, vol. 40, pp. 2112-5, 2015.
- [11] D. G. Matei, T. Legero, S. Häfner, C. Grebing, R. Weyrich, L. Zhang, L. Sonderhouse, J. M. Robinson, J. Ye, F. Riehle and U. Sterr, “1.5 μm lasers with sub-10 mHz linewidth,” *Phys Rev Lett*, p. 263202, 2017.
- [12] K. Bongs, Y. Singh, L. Smith, W. He, O. Kock, D. Swierad, J. Hughes, S. Schiller, S. Alighanbari, S. Origlia, S. Vogt, U. Sterr, C. Lisdat, R. Le Targat, J. Lodewyck, D. Holleville, B. Venon, S. Bize, G. P. Barwood, P. Gill, I. R. Hill, Y. B. Ovchinnikov, N. Poli, G. M. Tino, J. Stuhler and W. Kaenders, “Development of a strontium optical lattice clock for the SOC mission on the ISS,” *Comptes Rendus Physique*, vol. 16, pp. 553-64, 2015.
- [13] G. D. Cole, W. Zhang, M. J. Martin, J. Ye and M. Aspelmeyer, “Tenfold reduction of Brownian noise in high-reflectivity optical coatings,” *Nature Photonics*, vol. 7, p. 644, 2013.
- [14] A. Jiménez, T. Milde, N. Staacke, C. Aßmann, G. Carpintero and J. Sacher, “Narrow-line external cavity diode laser micro packaging in the NIR and MIR spectral range,” *Applied Physics B*, vol. 123, p. 207, 2017.
- [15] S. Webster, L. Wright and P. Gill, “Design Analysis of SRCs TN3: Finite-element thermo-mechanical analysis of SRC designs considered in WP2,” 2010.
- [16] P. D. Gaynor, I. R. Hill, R. J. Hendricks, G. P. Barwood and P. Gill, “Testbed Technical Note TN01b: Preliminary Design (TestBed-NPL-TN01b),” 2020.
- [17] T. Legero, T. Kessler and U. Sterr, “Tuning the thermal expansion properties of optical reference cavities with fused silica mirrors,” *J Opt Soc Am B*, vol. 27, pp. 914-9, 2010.
- [18] K. Nicklaus, M. Herding, X. Wang, N. Beller, O. Fitzau, M. Giesberts, M. Herper, G. P. Barwood, R. A. Williams, P. Gill, H. Koegel, S. A. Webster and M. Gohlke, “High Stability Laser For Next Generation Gravity Missions,” *ICSO Conference Proceedings 2014, Proc SPIE*, vol. 10563, pp. 105632T-2, 2017.
- [19] K. Dahl, K. Nicklaus, M. Herding, X. Wang, N. Beller, O. Fitzau, M. Giesberts, M. Herper, R. A. Williams, G. P. Barwood, P. Lovelock, P. Gill, H. Koegel and A. Sell, “High Stability Laser for Interferometric Earth Gravity Measurements,” *ICSO Conference Proceedings 2016, Proc SPIE*, vol. 10562, pp. 105620-J, 2017.
- [20] ESA, “Statement of Work: Multi-wavelength clock control unit MW-CCU (ESA-GSTP-TECMME-SOW-018014),” 2020.
- [21] “ORCS TN5: Performance evaluation and environmental test results,” 2016.
- [22] “OSRC Outcome test campaign 1064 nm system - Preliminary Report,” 2019.
- [23] F. B. Langitan and B. R. Lawn, “Hertzian Fracture Experiments on Abraded Glass Surfaces as Definitive Evidence for an Energy Balance Explanation of Auerbach's Law,” *Journal of Applied Physics*, vol. 40, no. 10, pp. 4009-4017, 1969.
- [24] ADS, “Development of an Acceleration Insensitive, Thermal Noise Mitigated, Optical Stabilising Reference Cavity (OSRC) Engineering Model: Volume 2 - Technical Proposal,” 2016.

- [25] A. Negrea and M. V. Predoi, “The elastic contact of a sphere with an elastic half-space, a comparison between analytical and finite elements solutions,” *UPB Scientific Bulletin, Series A: Applied Mathematics and Physics*, vol. 74, no. 4, pp. 69-78, 2012.
- [26] ADS, “Test plan cavity 1 (TN3b) Development of an acceleration insensitive, thermal noise mitigated, Optical Stabilising Reference Cavity (OSRC) Engineering Model,” 2018.
- [27] ESA, “Structural factors of safety for spaceflight hardware (ECSS-E-ST-32-10C Rev.2 Corr. 1),” 2019.
- [28] NPL, STI, “LSU Vibrational and Shock Test Plans (HSL2-NPL-TN-002b),” 2017.
- [29] ADS, “Performance evaluation cavity 1 (TN4) Development of an acceleration insensitive, thermal noise mitigated, Optical Stabilising Reference Cavity (OSRC) Engineering Model,” 2020.
- [30] STI, “OSRC Vibration Sensitivity Technology Review (TN1),” 2017.
- [31] F. C. Frank and B. R. Lawn, “On the theory of Hertzian fracture,” *Proceedings of the Royal Society*, vol. 299, pp. 291-306, 1969.
- [32] P. D. Warren, “Determining the Fracture Toughness of Brittle Materials by Hertzian Indentation,” *Journal of the European Ceramic Society*, vol. 15, pp. 201-207, 1995.
- [33] Corning Inc., “ULE 7972 Product Information Sheet,” 2016. [Online]. Available: <https://www.corning.com/media/worldwide/csm/documents/7972%20ULE%20Product%20Information%20Jan%202016.pdf>.
- [34] NPL, “TN01b: Preliminary Design (Optical Stabilising Reference Cavity – Optical and Optoelectronic Test Bed Enhancement activities (Test Bed)),” 2020.
- [35] L. Chen, J. L. Hall, J. Ye, T. Yang, E. Zang and T. Li, “Vibration-induced elastic deformation of Fabry-Perot cavities,” *Physical Review A*, vol. 74, 2006.
- [36] K. Numata, A. Kemery and J. Camp, “Thermal-Noise Limit in the Frequency Stabilization of Lasers with Rigid Cavities,” *Physical review letters*, 2004.
- [38] S. Webster and P. Gill, “Force-insensitive optical cavity,” *Optics Letters*, vol. 36, no. 18, 2011.
- [39] I. R. Hill, R. J. Hendricks, S. Donnellan, P. Gaynor, B. Allen, G. P. Barwood and P. Gill, “Dual-axis cubic cavity for drift-compensated multi-wavelength laser stabilisation,” *Optics Express*, vol. 29, p. 36758, 2021.
- [40] J. I. Thorpe, K. Numata and J. Livas, “Laser frequency stabilization and control through offset sideband locking to optical cavities,” *Optics Express*, vol. 16, p. 15980–15990, 2008.
- [41] J. Ye, L.-S. Ma and J. L. Hall, “Ultrasensitive detections in atomic and molecular physics: demonstration in molecular overtone spectroscopy,” *J. Opt. Soc. Am. B*, vol. 15, pp. 6-15, 1998.
- [42] L. Rothman et al, “The HITRAN2016molecularspectroscopicdatabase,” *Journal of Quantitative Spectroscopy & Radiative Transfer*, vol. 203, pp. 3-69, 2017.
- [43] R. Hobson, W. Bowden, A. Vianello, A. Silva, C. F. A. Baynham, H. S. Margolis, P. E. G. Baird, P. Gill and I. R. Hill, “A strontium optical lattice clock with 1×10^{-17} uncertainty and measurement of its absolute frequency,” *Metrologia*, vol. 57, p. 065026, 2020.
- [44] H. Kogelnik and T. Li, “Laser beams and resonators,” *Applied Optics*, vol. 5, pp. 1550-67, 1966.
- [45] W. Zhang, M. J. Martin, C. Benko, J. L. Hall, J. Ye, C. Hagemann, T. Legero, U. Sterr, F. Riehle, G. D. Cole and M. Aspelmeyer, “Reduction of residual amplitude modulation to 1×10^{-6} for frequency modulation and laser stabilization,” *Optics Letters*, vol. 39, pp. 1980-3, 2014.
- [46] L.-S. Ma, P. Jungner, J. Ye and J. L. Hall, “Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path,” *Optics Letters*, vol. 19, pp. 1777-9, 1994.
- [47] J. Van de Vegte, *Feedback Control Systems*, Prentice-Hall, Inc., 1986.
- [48] J. H. D.R.Leibrandt, “An open source digital servo for atomic, molecular, and optical physics experiments,” *Rev. Sci. Instrum.*, vol. 86, no. 12, pp. 1-12, 2015.

- [49] K. G. Libbrecht and J. L. Hall, “A lownoise highspeed diode laser current controller,” *Rev Sci Instrum*, vol. 64, pp. 2133-5, 1993.
- [50] M. Schioppo, J. Kronjäger, A. Silva, R. Ilieva, J. W. Paterson, C. F. A. Baynham, W. Bowden, I. R. Hill, R. Hobson, A. Vianello, M. Dovale-Álvarez, R. A. Williams, G. Marra, H. S. Margolis, A. Amy-Klein, O. Lopez, E. Cantin, H. Álvarez-Martínez, R. Le Targat, P. E. Pottie, N. Quintin, T. Legero, S. Häfner, U. Sterr, R. Schwarz, S. Dörscher, C. Lisdat, S. Koke, A. Kuhl, T. Waterholter, E. Benkler and G. Grosche, “Comparing ultrastable lasers at 7×10^{17} fractional frequency instability through a 2220 km optical fibre network,” *Nature Communications*, vol. 13, p. 212, 2022.
- [51] K. H. Brown, C. Morrow, S. Durbin and A. Baca, “Guideline for Bolted Joint Design and Analysis: Version 1.0,” Sandia National Laboratories, Albuquerque, NM, 2008.
- [52] NPL, PTB, SYRTE, “TN3: Finite-element thermo-mechanical analysis of SRC designs considered in WP2,” 2010.